arXiv:1405.2363v1 [cs.SY] 9 May 2014

A sampling-based approach to scalable constraint satisfaction in linear sampled-data systems—Part I: Computation*

Shahab Kaynama, Jeremy H. Gillula, Claire J. Tomlin†

(Preprint Submitted for Publication)

Abstract

Sampled-data (SD) systems, which are composed of both discrete- and continuous-time components, are arguably one of the most common classes of cyberphysical systems in practice; most modern controllers are implemented on digital platforms while the plant dynamics that are being controlled evolve continuously in time. As with all cyberphysical systems, ensuring hard constraint satisfaction is key in the safe operation of SD systems. A powerful analytical tool for guaranteeing such constraint satisfaction is the viability kernel: the set of all initial conditions for which a safety-preserving control law (that is, a control law that satisfies all input and state constraints) exists. In this paper we present a novel sampling-based algorithm that tightly approximates the viability kernel for high-dimensional sampled-data linear time-invariant (LTI) systems. Unlike prior work in this area, our algorithm formally handles both the discrete and continuous characteristics of SD systems. We prove the correctness and convergence of our approximation technique, provide discussions on heuristic methods to optimally bias the sampling process, and demonstrate the results on a twelve-dimensional flight envelope protection problem.

1 Introduction

The mathematical guarantee of satisfaction of hard input and state constraints is an increasingly desirable property that every safety-critical cyberphysical system must implement. The subset of the state space for which this property holds is known as the viability kernel [2], or alternatively, in the infinite horizon case, the maximal controlled-invariant set [3]. Consequently, a tremendous amount of work in the recent literature has been focused on methods for computing the viability kernel.

Constrained sampled-data (SD) systems describe a large class of cyberphysical systems. In most practical settings the system evolves continuously in time but the state is measured only at discrete time instants [4]. Consequently, any admissible control policy is piecewise constant; the input can only be applied at the beginning of each sampling interval and is kept constant (under zero-order hold) until the next sampling time. Examples of SD systems include control of blood glucose levels in type-1 diabetes [5] where insulin could only be administered every 30 minutes during the clinical trials. Additionally, restrictions on sampling frequency is not always due to sensory limitations or the particulars of an application. For instance, in model predictive control (MPC) a higher sampling frequency results in a significantly larger online optimization problem over each prediction horizon—a known limiting factor in embedded control design, e.g. in the automotive industry. Furthermore, without additional rate constraints, a higher sampling frequency would require controllers with much higher bandwidth.

Discretizing the dynamics and naively designing control policies in discrete time ignores the behavior of the true system in between two sampling instants. This inter-sample behavior can be crucial in safety-critical systems where the state constraint is associated with "safety" of the system. At the same time, performing continuous-time safety analysis on the system (for example, using the traditional level-set technique [6]), which only requires a mild assumption of

*An earlier version of this paper was partially presented at the 17th Int'l Conference on Hybrid Systems: Computation and Control, April 15-17, 2014, Berlin, Germany [1].

†The authors are with Electrical Engineering & Computer Sciences, University of California at Berkeley, 337 Cory Hall, Berkeley, CA 94720, USA. {kaynama, jgillula, tomlin}@eecs.berkeley.edu

Lebesgue measurability of the control input, cannot provide guarantees about the behavior of the SD system where the input is restricted to draw from the class of piecewise constant signals [7]. All this warrants a technique that can formally handle SD systems.

We present a sampling-based approach that yields tight under- and over-approximations of the viability kernel for *high-dimensional SD* linear time-invariant (LTI) dynamics.

1.1 Related Work

The classical numerical schemes for approximating the viability kernel are those based on gridding the state space (and discretizing the dynamics) and appropriately evolving the constraints over this stationary grid. Such schemes include the level-set methods [6] and variants of Saint-Pierre's viability algorithm [8]. Despite their versatility in handling complex dynamics and sets, the applicability of these schemes has historically been limited to systems of dimensions less than five. Efforts to generalize such grid-based techniques to moderately dimensioned systems include structure decomposition [9, 10] and approximate dynamic programming [11].

Algorithms from within the MPC community have also emerged that enable the computation of the viability kernel for discrete-time LTI systems with polytopic constraints [3]. Due to the fact that these algorithms recursively compute the Minkowski sum, linear transformation, and intersection of polytopes, they can only be applied to low dimensional systems; the number of vertices of the resulting polytope grows rapidly with each subsequent Minkowski sum operation, while the intersection operation at each iteration requires a vertex to facet enumeration of polytopes—an operation that is known to be intractable in high dimensions [12]. In more general contexts (e.g., for continuous-time systems), an ellipsoidal approximation of the region of attraction of the MPC is computed as a (crude) representation of the viability kernel [13]. These, as well as other approximation techniques such as β-contractive polytopes [14], generally require existence of a stabilizing controller within the constraints—a requirement that may not always be readily satisfied.

For LTI systems with convex constraints, [15] (discrete-time) and [16] (discrete- and continuous-time) introduced efficient and scalable algorithms based on support vector representations and piecewise ellipsoidal sets to conservatively approximate the viability kernel. These algorithms follow the flow of the dynamics in the same spirit as Lagrangian techniques for maximal reachability such as [17–19].

Approximating the viability kernel can also be viewed as a search for an appropriate control Lyapunov function subject to additional input and state constraints. As such, in parallel to the above developments, for polynomial systems with semi-algebraic constraints, various sum-of-squares (SOS)[1] optimization-based techniques have been proposed [20–24] that either directly form a polynomial approximation of the viability kernel, or can be modified to do so. The resulting bilinear SOS program is solved either by alternating search (prone to local optima) or through convex relaxations (introduces additional conservatism). The degree of the SOS multipliers, which directly translates to the presumed degree of the polynomial that is to describe the kernel, is commonly kept low (e.g. quartic), striking a tradeoff between excessive conservatism and computational complexity. A related SOS-based technique is the recent method of occupation measures [25, 26] that, while scalable, can only *over*-approximate the desired set (which is insufficient for safety). Though the approximation recovers the true kernel in the limit, the error is not monotonically decreasing with the degree of the multipliers.

All of the above algorithms are designed for either discrete- or continuous-time dynamics, and very little effort has been dedicated to *sampled-data* system—precisely the systems that, due to their physically descriptive and realistic nature, stand to benefit the most from safety formalism. To the best of our knowledge, the only other scalable work on computing the viability kernel for SD systems is presented in [7], where a piecewise ellipsoidal algorithm is proposed for LTI dynamics with ellipsoidal constraints. Unfortunately due to the projections and cross-products involved in this algorithm, the quality of the approximation degenerates rapidly with the time horizon.

1.2 Summary of Contributions

We propose a sampling-based approach that directly handles SD systems, albeit under LTI dynamics and convex constraints. The algorithm provably under-approximates the viability kernel in a scalable and efficient manner. Since

[1] SOS is a relaxation of the original semidefinite program.

the proposed method deals with individual trajectories to infer the evolution of sets of initial conditions, it is not explicitly restricted by a particular shape of the constraints, although its computational complexity does depend on these shapes.

Our technique yields a "tight" approximation in the sense that the resulting set touches the boundary of the true viability kernel from the inside with arbitrary precision up to a numerical constant. The algorithm is sampling based, meaning that the points to be included as boundary points of the viability kernel are sampled according to a probability distribution. We provide a convergence proof that shows that such a sampling-based technique is superior to any alternative deterministic approach.

Section 2 formulates the problem we wish to address. Section 3 presents our main under-approximation algorithm, while Section 4 proves its correctness and convergence. Section 5 discusses the computational complexity of the technique and showcases, via an example, its scalability. To provide a measure of conservatism of our technique, we will also present an over-approximation algorithm in Section 6. This over-approximation algorithm is then utilized in Section 7 to help guide the under-approximation process. This section also describes a few heuristics to bias the random sampling of the state constraint so as to achieve superior convergence and accuracy properties. We study a 12D flight envelope protection problem in Section 8, before providing concluding remarks and future directions in Section 9.

2 Problem Formulation

Consider the LTI system

$$\dot{x}(t) = Ax(t) + Bu(t) \tag{1}$$

with state $x(t) \in \mathbb{R}^n$, control input $u(t) \in \mathcal{U}$, where $\mathcal{U}$ is a compact convex subset of $\mathbb{R}^m$. A and B are constant matrices of appropriate dimension. The state of the system is measured at every time instant $t_k := k\delta$ for $k \in \mathbb{Z}_{\geq 0}$ and fixed sampling interval $\delta \in \mathbb{R}_{>0}$. We are concerned with the evolution of the system over $\mathbb{T} := [0, \tau]$ with arbitrary, finite time horizon $\tau \in \mathbb{R}_{>0}$. We denote by $N_\delta := \lceil \tau/\delta \rceil$ the number of sampling intervals in $\mathbb{T}$. The input is applied at the beginning of each sampling interval and is kept constant until the next sampling instant. Thus the input signal draws from the set of piecewise constant functions

$$\mathscr{U}_{\mathbb{T}}^{\text{pwc}} := \{u \colon \mathbb{T} \to \mathbb{R}^m \text{ piecewise const.},\ u(t_k) \in \mathcal{U}\ \forall k,\ u(t) = u(t_k)\ \forall t \in [t_k, t_{k+1})\}. \tag{2}$$

For every $x_0 \in \mathbb{R}^n$ and $u(\cdot) \in \mathscr{U}_{\mathbb{T}}^{\text{pwc}}$, we denote the (unique) trajectory of (1) by $x_{x_0}^u \colon \mathbb{T} \to \mathbb{R}^n$ with initial condition $x_{x_0}^u(0) = x_0$. When clear from the context, we shall drop the subscript and superscript from the trajectory notation.

For a nonempty compact convex state constraint set $\mathcal{K} \subset \mathbb{R}^n$, deemed *safe*, we examine the following construct:

Definition 1 (SD Viability Kernel)**.** *The finite-horizon SD viability kernel of $\mathcal{K}$ is the set of all initial states for which there exists a control law such that the trajectories emanating from those states remain in $\mathcal{K}$ over $\mathbb{T}$:*

$$\text{Viab}_{\mathbb{T}}^{\text{sd}}(\mathcal{K}) := \left\{x_0 \in \mathcal{K} \mid \exists u(\cdot) \in \mathscr{U}_{\mathbb{T}}^{\text{pwc}},\ \forall t \in \mathbb{T},\ x_{x_0}^u(t) \in \mathcal{K}\right\}. \tag{3}$$

We seek to find a scalable technique that *under-approximates* the viability kernel. (Any approximation must be conservative, in that at the very least all states for which no admissible input can maintain safety must be excluded.)

2.1 Preliminaries and Notation

We say that the vector field in (1) is bounded on $\mathcal{K}$ if $\exists M > 0$ such that $\|Ax + Bu\|_p \leq M\ \forall (x, u) \in \mathcal{K} \times \mathcal{U}$ for some norm $\|\cdot\|_p$. If $\mathcal{K}$ and $\mathcal{U}$ are compact, every continuous vector field is bounded on $\mathcal{K}$. The $\|\cdot\|_p$-distance of a point $x \in \mathbb{R}^n$ from a nonempty set $\mathcal{A} \subset \mathbb{R}^n$ is defined as $\text{dist}_p(x, \mathcal{A}) := \inf_{a \in \mathcal{A}} \|x - a\|_p$.

The *Minkowski sum* of any two nonempty subsets $\mathcal{A}$ and $\mathcal{C}$ is $\mathcal{A} \oplus \mathcal{C} := \{a + c \mid a \in \mathcal{A},\ c \in \mathcal{C}\}$; their *Pontryagin difference* (or, the *erosion* of $\mathcal{A}$ by $\mathcal{C}$) is $\mathcal{A} \ominus \mathcal{C} := \{a \mid a \oplus \mathcal{C} \subseteq \mathcal{A}\}$. We denote by $\partial\mathcal{C}$ the boundary of the set $\mathcal{C}$, by $\mathcal{C}^c$ its complement, and by $\text{vol}(\mathcal{C})$ its volume. $\text{conv}(\{v_0, \ldots, v_N\})$ denotes the convex hull of a set of points $v_0, \ldots, v_N$. $\mathcal{B}_p^n(x, a)$ denotes the closed p-norm ball of radius $a \in \mathbb{R}_{\geq 0}$ about a point x in $\mathbb{R}^n$, and $\mathcal{S}_p^{n-1}$ the codimension one boundary $\partial\mathcal{B}_p^n(0, 1)$. A *ray* in $\mathbb{R}^n$ is the set of points $\vec{r} = \{r_0 + sr_d \mid s \in \mathbb{R}_{\geq 0}\}$, where $r_0 \in \mathbb{R}^n$ is the *origin* of the ray, and $r_d \in \mathbb{R}^n$ is a unit vector giving the *direction* of the ray.

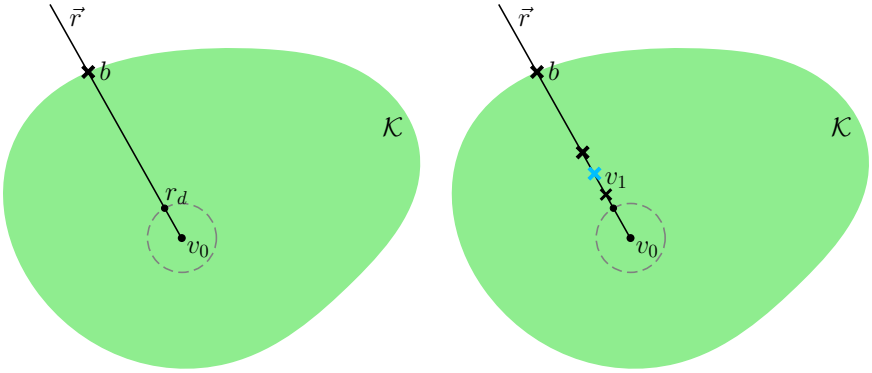

(a) Sample a point $r_d \in \mathcal{S}_2^{n-1}$ and form $\vec{r}$. Find $b = \vec{r} \cap \partial\mathcal{K}$.

(b) Perform bisection search on $\{v_0, b\}$ to find $v_1 \in \text{Viab}_{\mathbb{T}}^{\text{sd}}(\mathcal{K})$.

(c) Repeat for N samples.

(d) Return $\text{conv}(\{v_i\}_{i=0}^{N})$.

Figure 1: Illustration of Algorithm 1.

3 Methodology

3.1 Overview of the Algorithm

Before we describe our algorithm we first elaborate on some simple subroutines which we make use of, but which we will not formally define due to space constraints.

- FIND-INTERSECTION-ON-BOUNDARY$(\mathcal{C}, \vec{r})$ – Input: A convex compact set $\mathcal{C} \subset \mathbb{R}^n$, and a ray $\vec{r}$ with origin $r_0 \in \mathcal{C}$. Returns a point x along the ray $\vec{r}$ on $\partial\mathcal{C}$. We note that since $\mathcal{C}$ is convex and compact, and since $r_0 \in \mathcal{C}$, there is exactly one such point x. Runs in time proportional to the number of faces of $\mathcal{C}$ if $\mathcal{C}$ is a polytope, and constant time if $\mathcal{C}$ is an ellipsoid.
- SAMPLE-RAY(x) – Input: A point $x \in \mathbb{R}^n$. First samples a point r_d at random from $\mathcal{S}_2^{n-1}$. Returns the ray $\vec{r} = \{x + sr_d\}$. Runs in time linear in n.
- FEASIBLE$(x, \mathcal{C})$ – Input: A point $x \in \mathbb{R}^n$ and a convex compact set $\mathcal{C} \subset \mathbb{R}^n$. Returns true if $x \in \text{Viab}_{\mathbb{T}}^{\text{sd}}(\mathcal{C})$, and false otherwise. Further details on the implementation of this subroutine and its time complexity are given in Sections 3.2 and 5.

The under-approximation of the viability kernel proceeds as follows. We assume as input a description of $\mathcal{K}$,[2] as well as some initial point $v_0 \in \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$.[3] We then construct a polytopic approximation of $\mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$ by iteratively sampling a direction r_d and generating a ray $\vec{r}$ centered at v_0 (Algorithm 1, step 4); finding the point $b \in \mathbb{R}^n$ where $\vec{r}$ intersects the boundary of $\mathcal{K}$ (Algorithm 1, step 5); and then performing a bisection search along the line segment $\{v_0, b\}$ until we find a point v_i such that $v_i \in \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$ and $\mathrm{dist}_p(v_i, \partial\,\mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})) < \epsilon$ for some desired accuracy $\epsilon > 0$ in some norm $\|\cdot\|_p$ (Algorithm 2). By repeating for N samples and taking the convex hull of the resulting points $\{v_0, \ldots, v_N\}$ we arrive at a polytope $\mathcal{V}_N \subseteq \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$ which converges to $\mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$ (in a manner we shall formalize in Section 4.2). The conservatism of the algorithm will be formally proven in Section 4.1. Fig. 1 illustrates the under-approximation procedure.

Algorithm 1 Computes a polytopic under-approximation of $\mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$ with at most $N+1$ vertices

1: **procedure** POLYTOPIC-APPROX($\mathcal{K}, v_0, N$)
2: $\quad\mathcal{V}_0 \leftarrow \{v_0\}$
3: $\quad$**for** $i = 1$ to N **do** $\quad\triangleright\ N$ samples
4: $\quad\quad\vec{r} \leftarrow$ SAMPLE-RAY(v_0)
5: $\quad\quad b \leftarrow$ FIND-INTERSECTION-ON-BOUNDARY($\mathcal{K}, \vec{r}$)
6: $\quad\quad v_i \leftarrow$ BISECTION-FEASIBILITY($v_0, b, \mathcal{K}$)
7: $\quad\quad\mathcal{V}_i \leftarrow \mathrm{conv}(\mathcal{V}_{i-1} \cup \{v_i\})$
8: $\quad$**end for**
9: $\quad$**return** $\mathcal{V}_N$
10: **end procedure**

Algorithm 2 Determines an ϵ-accurate intersection of $\partial\,\mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$ and the line between a and b.

1: **function** BISECTION-FEASIBILITY($a, b, \mathcal{K}$)
2: $\quad c \leftarrow a + (b-a)/2$
3: $\quad$**if** FEASIBLE($c, \mathcal{K}$) **then**
4: $\quad\quad$**if** $\mathrm{dist}_p(b, c) < \epsilon$ **then**
5: $\quad\quad\quad$**return** c
6: $\quad\quad$**else**
7: $\quad\quad\quad$**return** BISECTION-FEASIBILITY($c, b, \mathcal{K}$)
8: $\quad\quad$**end if**
9: $\quad$**else**
10: $\quad\quad$**return** BISECTION-FEASIBILITY($a, c, \mathcal{K}$)
11: $\quad$**end if**
12: **end function**

3.2 Checking Point Feasibility

The key to our approach is the subroutine FEASIBLE($x_0, \mathcal{K}$) in Step 3 of Algorithm 2, which returns true only if $x_0 \in \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$. An overview of the procedure is as follows: Given an initial condition x_0, we verify the existence of a piecewise constant control law that ensures that the trajectory starting from x_0 belongs to a subset of $\mathcal{K}$ at every sampling instant. This subset is appropriately chosen at a certain distance from the exterior $\mathcal{K}^c$ such that the inter-sampling portions of the trajectory do not escape $\mathcal{K}$. The state x_0 is then labeled as *feasible*. We employ forward simulation to determine feasibility. To do this in a tractable fashion, we use a finite-difference approximation of

[2] In theory, $\mathcal{K}$ can be of any arbitrary (but convex) shape. In practice, this shape directly affects the run time complexity of the subroutines FIND-INTERSECTION-ON-BOUNDARY and FEASIBLE as discussed in the bullet points above.

[3] For linear systems if $(0,0) \in \mathcal{K} \times \mathcal{U}$ then $0 \in \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$, and thus we can use $v_0 = 0$. Otherwise, we assume we can sample points at random over $\mathcal{K}$ until we find a point x via the subroutine FEASIBLE($x, \mathcal{K}$) such that $x \in \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$.

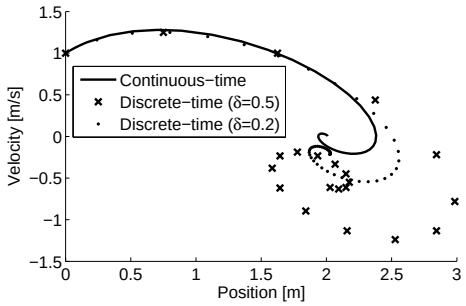

Figure 2: Simulated trajectory of a mass-spring-damper (unit parameters) under constant force in continuous time vs. its Euler discretization for two sampling times $\delta = 0.2, 0.5$.

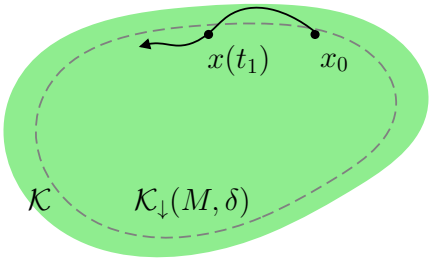

Figure 3: Erosion of $\mathcal{K}$ ensures that the curvature of the continuous trajectory in between two sampling instants cannot escape safety.

the dynamics. Therefore, we also need to take into account the effect of the discretization error (Fig. 2 depicts the significance of this error via a trivial example).

3.2.1 Dealing With Inter-Sampling Behavior

The following lemma ensures that if a control law exists that can keep the trajectory value evaluated at every sampling instant in a certain subset of $\mathcal{K}$, then the continuous evolution of the system in between sampling instants also maintains safety; that is, the curvature of the trajectory during the sampling intervals does not escape $\mathcal{K}$ (Fig. 3).

Lemma 1. *Let* (1) *be uniformly bounded on* $\mathcal{K}$ *by* $M > 0$ *in some norm* $\|\cdot\|_{p_1}$. *With a sampling interval* $\delta > 0$ *define*

$$\begin{aligned}\mathcal{K}_{\downarrow}(M,\delta) &:= \{x \in \mathcal{K} \mid \mathrm{dist}_{p_1}(x, \mathcal{K}^c) \geq M\delta\} \\ &= \mathcal{K} \ominus \mathcal{B}^n_{p_1}(0, M\delta),\end{aligned} \tag{4}$$

For a given initial condition x_0, *if* $\exists u(\cdot) \in \mathscr{U}_{\mathbb{T}}^{\mathrm{pwc}}$ *such that* $x(t_k) \in \mathcal{K}_{\downarrow}(M,\delta)\ \forall k \in \{0,\ldots,N_\delta\}$, *then* $x_0 \in \mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}(\mathcal{K})$.

The proof is similar to the continuous-time analysis in [16, Proof of Prop. 1] and is omitted here; see [1, Lem. 1] for more detail.

3.2.2 Dealing With Discretization Error

Recall that the solution $x_{k+1} = x(t_{k+1})$ of the SD system (1) at time t_{k+1} for any sampling interval $[t_k, t_k + \delta]$ starting from $x_k = x(t_k)$ using a constant input u_k is

$$x_{k+1} = e^{A\delta}x_k + \left(\int_0^\delta e^{A\lambda}d\lambda \cdot B\right)u_k. \tag{5}$$

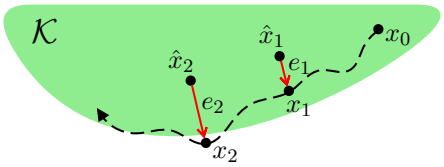

Figure 4: The mismatch between the true state x_k at sampling time t_k and the nominal state $\hat{x}_k$ of the discretized model. The error e_k propagates in time. If not accounted for in forward simulations, this mismatch could jeopardize safety.

Consider the Taylor series expansion $e^{As} = \sum_{i=0}^{\infty}(As)^i/i!$. To approximate the evolution of (1) at every discrete time instant t_k using a finite-difference equation, we approximate the above infinite sum by a ζth order finite sum

$$A_{\zeta,s} := \sum_{i=0}^{\zeta} \frac{(As)^i}{i!} \approx e^{As}, \quad \zeta < \infty \tag{6}$$

with truncation error

$$E_{\zeta,s} := e^{As} - A_{\zeta,s} = \sum_{i=\zeta+1}^{\infty} \frac{(As)^i}{i!}. \tag{7}$$

The trajectory of the resulting finite-difference equation is

$$\hat{x}_{k+1} = A_{\zeta,\delta}\hat{x}_k + \left(\int_0^{\delta} A_{\zeta,\lambda} d\lambda \cdot B\right) u_k. \tag{8}$$

We refer to $\hat{x}$ as the *nominal* state of the system.

It is important to emphasize that the computation of the matrix exponential has always been a challenging task [27]. An exact computation is generally not possible. Instead, approximate techniques are employed. For instance, Matlab uses the Padé approximation with scaling and squaring method of [28]. In most practical cases, the forward Euler approximation ($\zeta = 1$) is used.

Truncating the tail E of the Taylor series expansion introduces a discretization error that results in mismatch between the true values of the system trajectory at discrete time instants and the values generated by the finite-difference model. That is, the *discretization error* e_k at time t_k resulting from the truncation error E will cause a deviation from the true state $x_k = x(t_k)$ such that

$$\hat{x}_k = x_k - e_k. \tag{9}$$

Consequently, to formally guarantee an under-approximation of the SD viability kernel using a time discretized approach—i.e. via simulation of the nominal trajectory—we must take into account the effect of the error e and its forward propagation in time (Fig. 4).

Lemma 2. *Suppose that there exist constants $\gamma_k \geq 0$ and a norm $\|\cdot\|_{p_2}$ such that $\|e_k\|_{p_2} \leq \gamma_k \; \forall k$. Let*

$$\mathcal{K}^k_{\downarrow}(M, \delta, \gamma_k) := \mathcal{K}_{\downarrow}(M, \delta) \ominus \mathcal{B}^n_{p_2}(0, \gamma_k) \neq \emptyset. \tag{10}$$

Then we have that

$$\hat{x}_k \in \mathcal{K}^k_{\downarrow}(M, \delta, \gamma_k) \Rightarrow x_k \in \mathcal{K}_{\downarrow}(M, \delta). \tag{11}$$

Proof. Regardless of the perturbation caused by the error, since $e_k \in \mathcal{B}^n_{p_2}(0, \gamma_k)$ we have that $x_k \in \{\hat{x}_k\} \oplus \mathcal{B}^n_{p_2}(0, \gamma_k)$. By enforcing the condition $\hat{x}_k \in \mathcal{K}^k_{\downarrow}(M, \delta, \gamma_k)$ we guarantee that $x_k \in \left(\mathcal{K}_{\downarrow}(M, \delta) \ominus \mathcal{B}^n_{p_2}(0, \gamma_k)\right) \oplus \mathcal{B}^n_{p_2}(0, \gamma_k) \subseteq \mathcal{K}_{\downarrow}(M, \delta)$ since for any nonempty sets $\mathcal{X}$ and $\mathcal{Y}$ it holds true that $(\mathcal{X} \ominus \mathcal{Y}) \oplus \mathcal{Y} \subseteq \mathcal{X}$. □

For any given initial condition x_0 we trivially have $e_0 = 0$ and thus $\mathcal{K}^0_{\downarrow}(M, \delta, \gamma_0) = \mathcal{K}_{\downarrow}(M, \delta)$. We now describe a procedure to compute error bounds γ_k which will be used for *a priori* construction of the sets $\mathcal{K}^k_{\downarrow}(M, \delta, \gamma_k)$.

Using the identity $e^{As} = A_{\zeta,s} + E_{\zeta,s}$ in (5) yields

$$x_k = \underbrace{A_{\zeta,\delta} x_{k-1} + \left(\int_0^\delta A_{\zeta,\lambda} d\lambda \cdot B\right) u_{k-1}}_{\hat{x}_k} + \underbrace{E_{\zeta,\delta} x_{k-1} + \left(\int_0^\delta E_{\zeta,\lambda} d\lambda \cdot B\right) u_{k-1}}_{e_k}. \tag{12}$$

Back-substituting the solutions x_{k-1} into x_k, and $\hat{x}_{k-1}$ into $\hat{x}_k$ for every $k = 1, \ldots, N_\delta$ we can rewrite $e_k = x_k - \hat{x}_k$ as

$$e_k = \left((A_{\zeta,\delta} + E_{\zeta,\delta})^k - A_{\zeta,\delta}^k\right) x_0 + \sum_{i=0}^{k-1} \Bigg[\left((A_{\zeta,\delta} + E_{\zeta,\delta})^i - A_{\zeta,\delta}^i\right) \times \int_0^\delta A_{\zeta,\lambda} d\lambda + (A_{\zeta,\delta} + E_{\zeta,\delta})^i \int_0^\delta E_{\zeta,\lambda} d\lambda \Bigg] B u_{k-1-i}. \tag{13}$$

We can do so because for any discrete-time LTI system $x_{k+1} = \Phi x_k + \Psi u_k$, the solution x_k can be written in terms of the initial condition x_0 and the past inputs as $x_k = \Phi^k x_0 + \sum_{i=0}^{k-1} \Phi^i \Psi u_{k-1-i}$.

To compute the upper-bound γ_k on (13), invoke i) the inequality $\left\|A^k\right\| \leq \|A\|^k$ which holds for any matrix A and positive constant k, ii) the multiplicative and triangular inequalities, and iii) the binomial expansion for $(A_{\zeta,\delta} + E_{\zeta,\delta})^k$ (which is valid since $A_{\zeta,s} E_{\zeta,s} = E_{\zeta,s} A_{\zeta,s}$ [29]):

$$\begin{aligned} \|e_k\| \leq & \left\| \sum_{l=0}^{k} \binom{k}{l} A_{\zeta,\delta}^l E_{\zeta,\delta}^{k-l} - A_{\zeta,\delta}^k \right\| \|x_0\| \\ & + \sum_{i=0}^{k-1} \left[\left\| \sum_{l=0}^{i} \binom{i}{l} A_{\zeta,\delta}^l E_{\zeta,\delta}^{i-l} - A_{\zeta,\delta}^i \right\| \left\| \int_0^\delta A_{\zeta,\lambda} d\lambda \right\| \right. \\ & \left. + \|A_{\zeta,\delta} + E_{\zeta,\delta}\|^i \left\| \int_0^\delta E_{\zeta,\lambda} d\lambda \right\| \right] \sup_{u \in \mathcal{U}} \|Bu\|. \end{aligned} \tag{14}$$

For $k = 1$ this inequality is

$$\|e_1\| \leq \|E_{\zeta,\delta}\| \|x_0\| + \left\| \int_0^\delta E_{\zeta,\lambda} d\lambda \right\| \sup_{u \in \mathcal{U}} \|Bu\|. \tag{15}$$

For $k > 1$ we find

$$\begin{aligned} \|e_k\| \leq & \sum_{l=0}^{k-1} \binom{k}{l} \left\|A_{\zeta,\delta}^l\right\| \|E_{\zeta,\delta}\|^{k-l} \|x_0\| \\ & + \sum_{i=1}^{k-1} \left[\sum_{l=0}^{i-1} \binom{i}{l} \left\|A_{\zeta,\delta}^l\right\| \|E_{\zeta,\delta}\|^{i-l} \left\| \int_0^\delta A_{\zeta,\lambda} d\lambda \right\| \right. \\ & \left. + \left(\|A_{\zeta,\delta}\| + \|E_{\zeta,\delta}\|\right)^i \left\| \int_0^\delta E_{\zeta,\lambda} d\lambda \right\| \right] \sup_{u \in \mathcal{U}} \|Bu\|. \end{aligned} \tag{16}$$

The term $\int_0^\delta A_{\zeta,\lambda} d\lambda$ in (16) is a definite integral of a finite sum and it can be easily computed:

$$\int_0^\delta A_{\zeta,\lambda} d\lambda = \sum_{i=0}^{\zeta} \int_0^\delta \frac{(A\lambda)^i}{i!} d\lambda = \sum_{i=0}^{\zeta} \frac{A^i \delta^{i+1}}{(i+1)!}. \tag{17}$$

Evaluating the integral $\int_0^\delta E_{\zeta,\lambda} d\lambda$ is trickier. We can, however, compute an upper-bound on its ∞-norm. We will use the following property [30]:

$$\|E_{\zeta,\delta}\|_\infty \leq \frac{(\|A\|_\infty \delta)^{\zeta+1}}{(\zeta+1)!} \cdot \frac{1}{1-\varepsilon} =: \psi_\delta, \tag{18}$$

where the discretization order ζ is chosen such that $\varepsilon := \frac{\|A\|_\infty \delta}{\zeta+2} < 1$ (to ensure convergence of the power series $1+\varepsilon+\varepsilon^2+\cdots$). Similarly, we can derive

$$\begin{aligned}\left\|\int_0^\delta E_{\zeta,\lambda} d\lambda\right\|_\infty &\leq \int_0^\delta \|E_{\zeta,\lambda}\|_\infty d\lambda \leq \int_0^\delta \sum_{i=\zeta+1}^{\infty} \frac{\|A\|_\infty^i \lambda^i}{i!} d\lambda \\ &= \sum_{i=\zeta+1}^{\infty} \int_0^\delta \frac{\|A\|_\infty^i \lambda^i}{i!} d\lambda = \sum_{i=\zeta+1}^{\infty} \frac{\|A\|_\infty^i \delta^{i+1}}{(i+1)!} \\ &\leq \frac{\|A\|_\infty^{\zeta+1} \delta^{\zeta+2}}{(\zeta+2)!} \cdot \frac{1}{1-\eta} \leq \psi_\delta \cdot \frac{\delta}{\zeta+2}, \end{aligned} \tag{19}$$

where $\eta = \varepsilon \cdot \left(1 - \frac{1}{\zeta+3}\right)$. Note that $\varepsilon < 1 \Rightarrow \eta < 1$, which automatically ensures the convergence of the power series $1+\eta+\eta^2+\cdots$ used in derivation of (19).

Substituting (17)–(19) into (15)–(16) for the ∞-norm we obtain a conservative bound $\tilde{\gamma}_k$ on $\|e_k\|_\infty$ as

$$\tilde{\gamma}_1 := \psi_\delta \|x_0\|_\infty + \left(\psi_\delta \cdot \tfrac{\delta}{\zeta+2}\right) \sup_{u\in\mathcal{U}} \|Bu\|_\infty ; \tag{20}$$

$$\begin{aligned}\tilde{\gamma}_k := &\sum_{l=0}^{k-1} \binom{k}{l} \left\|A_{\zeta,\delta}^l\right\|_\infty \psi_\delta^{k-l} \|x_0\|_\infty \\ &+ \sum_{i=1}^{k-1} \left[\sum_{l=0}^{i-1} \binom{i}{l} \left\|A_{\zeta,\delta}^l\right\|_\infty \psi_\delta^{i-l} \left\| \sum_{j=0}^{\zeta} \frac{A^j \delta^{j+1}}{(j+1)!} \right\|_\infty \right. \\ &\left. + \left(\|A_{\zeta,\delta}\|_\infty + \psi_\delta\right)^i \psi_\delta \cdot \tfrac{\delta}{\zeta+2} \right] \sup_{u\in\mathcal{U}} \|Bu\|_\infty \end{aligned} \tag{21}$$

for $k > 1$. Using the upper-bounds (20)–(21) in Lemma 2 allows us to check for feasibility of a given point x_0 via only the finite-difference model, while ensuring that safety will not be violated due to discretization.

The bound $\tilde{\gamma}_k$ is asymptotically tight in the sense that for any k, $\lim_{\zeta\to\infty} \tilde{\gamma}_k = 0$. In practice, the chosen order of discretization ζ must be large enough so as to ensure convergence of the power series in derivation of (18) as well as non-emptiness of the eroded sets in Lemma 2.

3.2.3 Verifying Feasibility of x_0 via Forward Simulation

We can now simply use the discretized model

$$\hat{x}_{k+1} = A_{\zeta,\delta}\hat{x}_k + B_{\zeta,\delta} u_k \tag{22}$$

with $B_{\zeta,\delta} := \int_0^\delta A_{\zeta,\lambda} d\lambda B$ to determine feasibility of a given initial condition x_0 without worrying about the discretization error or the inter-sampling behavior of the continuous trajectories of (1) and their potentially negative impact on safety: If there exists a sequence of controls $\{u_k\}$ so that the nominal states $\hat{x}_k$ of the closed-loop system belong to the precomputed sets $\mathcal{K}_\downarrow^k(M, \delta, \tilde{\gamma}_k)$ as described above, then via Lemmas 2 and 1 the trajectories of (1) never exit $\mathcal{K}$.

Let us now construct the prediction equation as in (23) (where we have used the notation $Gx_0 + H\mathbf{u}$ to abbreviate the right-hand side of the equality), and formulate the finite horizon feasibility program

$$\begin{bmatrix} \hat{x}_0 \\ \hat{x}_1 \\ \hat{x}_2 \\ \vdots \\ \hat{x}_{N_\delta} \end{bmatrix} = \underbrace{\begin{bmatrix} I \\ A_{\zeta,\delta} \\ A_{\zeta,\delta}^2 \\ \vdots \\ A_{\zeta,\delta}^{N_\delta} \end{bmatrix}}_{G} x_0 + \underbrace{\begin{bmatrix} 0 & 0 & \dots & 0 \\ B_{\zeta,\delta} & 0 & \dots & 0 \\ A_{\zeta,\delta}B_{\zeta,\delta} & B_{\zeta,\delta} & \dots & 0 \\ \vdots & \vdots & \dots & \vdots \\ A_{\zeta,\delta}^{N_\delta-1}B_{\zeta,\delta} & A_{\zeta,\delta}^{N_\delta-2}B_{\zeta,\delta} & \dots & B_{\zeta,\delta} \end{bmatrix}}_{H} \underbrace{\begin{bmatrix} u_0 \\ u_1 \\ \vdots \\ u_{N_\delta-1} \end{bmatrix}}_{\mathbf{u}} \tag{23}$$

$$\min_{\mathbf{u}} \quad 0 \tag{24a}$$

$$\text{subj. to} \quad \mathbf{u} \in \mathcal{U}^{N_\delta} \tag{24b}$$

$$\hat{x}_k \in \mathcal{K}_\downarrow^k(M, \delta, \tilde{\gamma}_k), \quad k = 0, \dots, N_\delta \tag{24c}$$

$$[\hat{x}_0 \ \cdots \ \hat{x}_{N_\delta}]^\top = Gx_0 + H\mathbf{u}. \tag{24d}$$

Theorem 1. *If* $\exists \mathbf{u}^*$ *satisfying* (24)*, then* $x_0 \in \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$.

Proof. The proof follows directly from Lemmas 1 and 2. More specifically, the prediction equation (24d), for a fixed input sequence $\mathbf{u}$, generates a forward simulation of the finite-difference model (22) over the desired horizon $[0, N_\delta] \cap \mathbb{Z}$ corresponding to the continuous time horizon $[0, N_\delta\delta] = \mathbb{T}$. Constraint (24b) ensures that this input sequence is pointwise admissible (meaning that every member of the sequence belongs to $\mathcal{U}$), while constraint (24c) restricts $\hat{x}_k$ so that the trajectory of (1) evaluated at t_k belongs to $\mathcal{K}_\downarrow(M, \delta)$ since via Lemma 2 $\hat{x}_k \in \mathcal{K}_\downarrow^k(M, \delta, \tilde{\gamma}_k) \Rightarrow x_k = x(t_k) \in \mathcal{K}_\downarrow(M, \delta)$. Lemma 1 then automatically guarantees that $x(t) \in \mathcal{K} \ \forall t \in \mathbb{T}$ which implies $x_0 \in \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$. □

The subroutine FEASIBLE in Algorithm 2 employs Theorem 1 to determine the feasibility of a given sample point x_0. Its computational complexity is proportional to the complexity of (24) which, with polytopic constraints for example, is simply a linear program (LP).[4]

4 Conservatism and Convergence

4.1 Algorithm Correctness

Theorem 2. *Given convex sets* $\mathcal{K}$ *and* $\mathcal{U}$ *and an initial point* $v_0 \in \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$*,* $\mathcal{V}_N = \text{POLYTOPIC-APPROX}(\mathcal{K}, v_0, N)$ *is a subset of* $\mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K}) \ \forall N$.

Proof. (By induction) First, it is obvious that for $N = 0$, $\mathcal{V}_0 = \text{POLYTOPIC-APPROX}(\mathcal{K}, v_0, 0) = \{v_0\} \subseteq \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$, since we are given that $v_0 \in \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$.

Next, assume that $\mathcal{V}_{N-1} = \text{POLYTOPIC-APPROX}(\mathcal{K}, v_0, N-1) = \mathrm{conv}(\{v_0, \dots, v_{N-1}\}) \subseteq \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$, and that we have $v_N = \text{BISECTION-FEASIBILITY}(v_0, b, \mathcal{K})$ for some point $b \in \partial\mathcal{K}$. Let $\mathcal{V}_N = \mathrm{conv}(\mathcal{V}_{N-1} \cup \{v_N\})$. Since BISECTION-FEASIBILITY only returns points which are inside $\mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$, we know $v_N \in \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$. Now since $\mathcal{V}_N$ is convex, $\forall x_0 \in \mathcal{V}_N \ \exists x_0' \in \mathcal{V}_{N-1}$ and $\exists \theta \in [0,1]$ s.t. $x_0 = \theta x_0' + (1-\theta)v_N$. For x_0' we know (by induction hypothesis) $\exists u_{x_0'}(\cdot) \in \mathscr{U}_{\mathbb{T}}^{\mathrm{pwc}}$ s.t. $x'(t) = e^{At}x_0' + \int_0^t e^{A(t-r)}Bu_{x_0'}(r)dr \in \mathcal{K} \ \forall t \in \mathbb{T}$. For v_N we also know $\exists u_{v_N}(\cdot) \in \mathscr{U}_{\mathbb{T}}^{\mathrm{pwc}}$ s.t. $x''(t) = e^{At}v_N + \int_0^t e^{A(t-r)}Bu_{v_N}(r)dr \in \mathcal{K} \quad \forall t \in \mathbb{T}$. Therefore,

$$\begin{aligned} \tilde{x}(t) := \theta x'(t) + (1-\theta)x''(t) &= e^{At}\left(\theta x_0' + (1-\theta)v_N\right) \\ &\quad + \int_0^t e^{A(t-r)}B\left(\theta u_{x_0'}(r) + (1-\theta)u_{v_N}(r)\right)dr \\ &= e^{At}x_0 + \int_0^t e^{A(t-r)}Bu_{x_0}(r)dr \in \mathcal{K} \ \ \forall t \in \mathbb{T} \end{aligned} \tag{25}$$

[4] We note that the vast majority of the calculations for $\tilde{\gamma}_k$ from (20)–(21) can be done once and ahead of time. Online, to be able to construct $\mathcal{K}_\downarrow^k$ in (24c), two simple operations (a multiplication by $\|x_0\|$ and an addition) are all that is needed to form $\tilde{\gamma}_k$.

since $\mathcal{K}$ and $\mathcal{U}$ are convex and compact. Thus, $u_{x_0}(\cdot) = \theta u_{x_0'}(\cdot) + (1-\theta)u_{v_N}(\cdot) \in \mathscr{U}_{\mathbb{T}}^{\text{pwc}}$ is safety-preserving and $x_0 \in \text{Viab}_{\mathbb{T}}^{\text{sd}}(\mathcal{K})$. Because x_0 was chosen arbitrarily in $\mathcal{V}_N$, we conclude that $\mathcal{V}_N \subseteq \text{Viab}_{\mathbb{T}}^{\text{sd}}(\mathcal{K})$. □

4.2 Algorithm Convergence

One of the striking features of our algorithm is that it is random; additional points on the boundary of the polytopic approximation of $\text{Viab}_{\mathbb{T}}^{\text{sd}}(\mathcal{K})$ are iteratively generated based on a random sampling. This random nature is due to the fact that $\partial\,\text{Viab}_{\mathbb{T}}^{\text{sd}}(\mathcal{K})$ is unknown *a priori* (and is, in fact, what we are trying to estimate) so it is impossible to know what points to sample to construct a polytope that converges to $\text{Viab}_{\mathbb{T}}^{\text{sd}}(\mathcal{K})$ as quickly as possible. In fact, any algorithm which deterministically chose points for which to verify feasibility could be presented with a safe set $\mathcal{K}$ and system dynamics for which the algorithm would converge arbitrarily poorly. This fact is related to results in the literature of estimating the volume of convex bodies using a *separation oracle*[5]. More specifically, part of the literature on algorithms for estimating the volume of convex bodies states that it can be shown that for any algorithm that deterministically queries a separation oracle a polynomial number of times to build a polytopic approximation, the error (the difference in volume between the approximation and the true set) could be exponential in the number of dimensions [31]. Random algorithms, on the other hand, can perform in a provably better manner [32].

To prove our algorithm's asymptotic convergence, first let $\overline{\text{Viab}_{\mathbb{T}}^{\text{sd}}}(\mathcal{K})$ be the subset of the viability kernel we are actually attempting to approximate, i.e.

$$\overline{\text{Viab}_{\mathbb{T}}^{\text{sd}}}(\mathcal{K}) := \{x_0 \mid \exists u(\cdot) \in \mathscr{U}_{\mathbb{T}}^{\text{pwc}},\, \forall k,\, x(t_k) \in \mathcal{K}_{\downarrow}(M,\delta)\}. \tag{26}$$

While the true kernel contains all initial conditions for which a piecewise constant control keeps $x(t) \in \mathcal{K}$, the above set only encompasses initial conditions for which $x(t_k) \in \mathcal{K}_{\downarrow}(M,\delta)$ (which is a sufficient, but not necessary, condition to imply $x(t) \in \mathcal{K}$; cf. Lemma 1). Define the volumetric error between these two sets as

$$\epsilon_{\text{cont}}(M\delta) := \text{vol}(\text{Viab}_{\mathbb{T}}^{\text{sd}}(\mathcal{K})) - \text{vol}(\overline{\text{Viab}_{\mathbb{T}}^{\text{sd}}}(\mathcal{K})) \tag{27}$$

and note that it depends only on the term $M\delta$, due to the definition of $\mathcal{K}_{\downarrow}(M,\delta)$ in (4).

Next, consider the output $\mathcal{V}_N$ of Algorithm 1. Clearly, the accuracy of under-approximation of the viability kernel by the set $\mathcal{V}_N$ is implicitly dependant on the discretization order ζ (Section 3.2.2), and the accuracy ϵ of the bisection search (Algorithm 2). To reflect this dependency, we adapt the extended notation $\mathcal{V}_N^{\zeta,\epsilon}$.[6] Evidently, for fixed values of ζ and ϵ as $N \to \infty$, this set only approximates a subset $\overline{\text{Viab}_{\mathbb{T}}^{\text{sd}}}(\mathcal{K},\zeta,\epsilon)$ of the set $\overline{\text{Viab}_{\mathbb{T}}^{\text{sd}}}(\mathcal{K})$:

$$\limsup_{N\to\infty}\ \mathcal{V}_N^{\zeta,\epsilon} = \overline{\text{Viab}_{\mathbb{T}}^{\text{sd}}}(\mathcal{K},\zeta,\epsilon) \tag{28}$$

$$\limsup_{\zeta\to\infty,\epsilon\to 0}\ \overline{\text{Viab}_{\mathbb{T}}^{\text{sd}}}(\mathcal{K},\zeta,\epsilon) = \overline{\text{Viab}_{\mathbb{T}}^{\text{sd}}}(\mathcal{K}) \tag{29}$$

with $\limsup$ denoting the Kuratowski upper-limit. We are now ready to present our algorithm's convergence property.

Proposition 1 (Rate of Convergence)**.** *Let $\epsilon_{\text{vol}}(N,\zeta,\epsilon)$ be the volumetric error between the viability kernel and the output of our algorithm, minus the error $\epsilon_{\text{cont}}(M\delta)$ between the true kernel and $\overline{\text{Viab}_{\mathbb{T}}^{\text{sd}}}(\mathcal{K})$:*

$$\epsilon_{\text{vol}}(N,\zeta,\epsilon) := \text{vol}(\text{Viab}_{\mathbb{T}}^{\text{sd}}(\mathcal{K})) - \text{vol}(\mathcal{V}_N^{\zeta,\epsilon}) - \epsilon_{\text{cont}}(M\delta).$$

Then our algorithm converges as

$$\lim_{\substack{N\to\infty\\ \zeta\to\infty\\ \epsilon\to 0}} \epsilon_{\text{vol}}(N,\zeta,\epsilon) N^{\frac{2}{n-1}} = c_n(\text{Viab}_{\mathbb{T}}^{\text{sd}}(\mathcal{K}), M\delta), \tag{30}$$

where $c_n(\text{Viab}_{\mathbb{T}}^{\text{sd}}(\mathcal{K}), M\delta)$ is a constant dependent on the dimension n, the shape of the viability kernel (more specifically its Gauss-Kronecker curvature), and the value $M\delta$.

[5] A separation oracle is a function that accepts as input a convex set and a point, and returns whether or not that point is inside the convex set. In our algorithm FEASIBLE$(x,\mathcal{K})$ plays this role.

[6] Theorem 2, restated in terms of the extended notation, asserts that $\mathcal{V}_N^{\zeta,\epsilon} \subseteq \text{Viab}_{\mathbb{T}}^{\text{sd}}(\mathcal{K})\ \forall \zeta,\epsilon,N$.

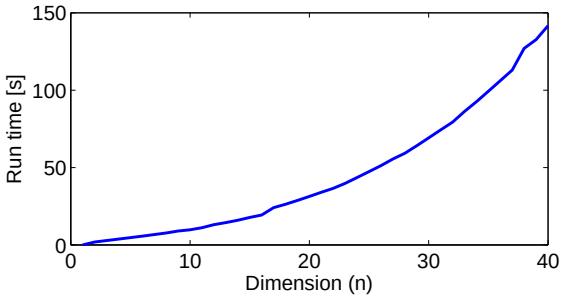

Figure 5: Run time of the algorithm for a chain of n integrators.

The proof requires some background on random algorithms for convex bodies and is provided in Appendix A.

Proposition 1 asserts that, for fixed dimension n, the volumetric error between the outcome of our algorithm and the true viability kernel asymptotically converges, at the exponential rate of $\hat{c}_n(\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}(\mathcal{K}))/N^{\frac{2}{n-1}}$, to a numerical constant due to the sampled-data nature of the system. On the other hand, to keep the accuracy of the approximation the same as $n \to kn$ we would need an increase of $N \to N^{\frac{kn-1}{n-1}}$. However, the fact that we only store samples on the boundary of the viability kernel to describe that set (as opposed to storing a grid of the entire set $\mathcal{K}$ and possibly beyond) requires significantly less memory than conventional approaches such as the SD level-set method in [7]. The flexibility in choosing the number of samples strikes a direct tradeoff between accuracy and computational complexity, making our algorithm scalable to high dimensions. The computed approximation is far more accurate (and quite possibly more scalable) than the piecewise ellipsoidal technique also presented in [7].

Again, due to the results in [33], the above convergence rate is optimal (up to a multiplicative constant depending on the probability density function used for sampling); no other algorithm that approximates the kernel by sampling from its boundary will be able to converge at a faster rate.

5 Computational Complexity & Scalability

The run time complexity of our algorithm (for fixed number of sampling intervals N_δ) is $\mathcal{O}(N \log(d)\Phi(n))$, where N is the number of samples/vertices, d is the "diameter" of the set $\mathcal{K}$, and Φ is the complexity of the feasibility program (24) as a function of the state dimension n. That is, the algorithm runs in time linear in the number of samples N, logarithmic in the diameter d of $\mathcal{K}$ due to complexity of the bisection search, and proportional to Φ in the complexity of the appropriate feasibility program (24). For instance, with polytopic constraints, the feasibility problem is an LP and thus the algorithm runs in time sub-cubic in n. This is a direct improvement over existing techniques for approximating the SD viability kernel. Furthermore, since each vertex is processed completely independently of others, our algorithms is highly parallelizable.

To demonstrate the scalability of our algorithm, consider the chain of n integrators $d^n x/dt^n = u$ with constraints $\mathcal{U} = [-0.15, 0.15]$ and $\mathcal{K} = \{x \mid \|x\|_\infty \leq 0.5\}$. The state is measured every $\delta = 0.05\,\mathrm{s}$ and safety is to be maintained over $\mathbb{T} = [0, 1]$. We use a discretization order of $\zeta = 4$, bisection accuracy of $\epsilon = 0.01$ with maximum of three-level bisection depth, and employ YALMIP [34] to implement (24) and MPT [35] for simple operations with polytopes. All of these parameters are kept constant as we increase the dimension n and the number of samples $N = 2n$ to examine the scalability of our algorithm. The results are shown in Fig. 5. The algorithm is implemented in MATLAB R2011b and tested on an Intel Core i7 at $2.9\,\mathrm{GHz}$ with $16\,\mathrm{GB}$ RAM running 64-bit Windows 7 Pro (without optimizing the code for speed).

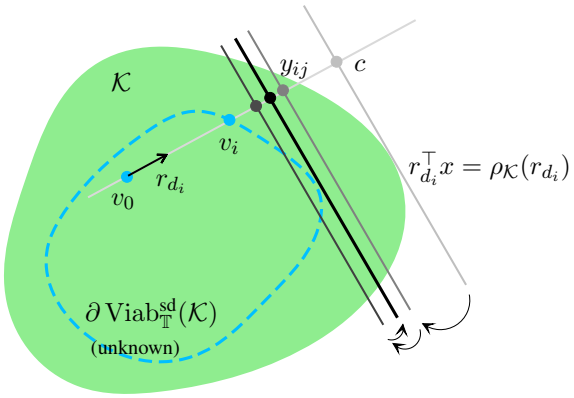

Figure 6: An ϵ_{o}-accurate bisection search between v_i and c determines via (33) the supporting hyperplane of the viability kernel.

6 Bounding the Error for Finite Number of Samples: Computing a Tight Over-Approximation

Every convex set can be over-approximated by any finite collection of its *support functions*. The support function of a convex compact set $\mathcal{C} \subset \mathbb{R}^n$ along $\ell \in \mathbb{R}^n$ is

$$\rho_{\mathcal{C}}(\ell) := \max_{x \in \mathcal{C}} \ell^\top x. \tag{31}$$

The half-space $\{x \mid \ell^\top x \leq \rho_{\mathcal{C}}(\ell)\}$ contains $\mathcal{C}$, and the hyper-plane $\{x \mid \ell^\top x = \rho_{\mathcal{C}}(\ell)\}$ is a supporting hyperplane for $\mathcal{C}$ with normal vector ℓ and distance value $\rho_{\mathcal{C}}(\ell)$. It follows that $\mathcal{C} \subseteq \bigcap_{\ell \in \mathcal{L}} \{x \mid \ell^\top x \leq \rho_{\mathcal{C}}(\ell)\}$ with $\mathcal{L}$ a finite subset of $\mathbb{R}^n$.

Let r_{d_i} be the direction r_d along which we have determined the vertex v_i of the under-approximation set through the ith iteration of Algorithm 1. It is easy to compute the support function of the set $\mathcal{K}$ along this direction: $\rho_{\mathcal{K}}(r_{d_i}) = \max_{x \in \mathcal{K}} r_{d_i}^\top x$. To find the supporting hyperplane of the true, unknown viability kernel (or rather some appropriate approximation of it) in the direction r_{d_i}, we move the hyperplane

$$\{x \mid r_{d_i}^\top x = \rho_{\mathcal{K}}(r_{d_i})\} \tag{32}$$

on the interior of $\mathcal{K}$ until we find at least one feasible point on this plane that belongs to the kernel (or its over-approximation); see Fig. 6.

We do so iteratively by first performing an ϵ_{o}-accurate bisection search between the point v_i and the point of intersection c of the ray $\vec{r}_i$ passing through v_i and v_0 with the hyperplane (32). This gives us points y_{ij} indexed by each step j of the new bisection search. We then solve the following modified convex feasibility program *for every* y_{ij} for the given direction r_{d_i}:

$$\begin{aligned}
\min_{\mathbf{u}, x_0} \quad & 0 && \text{(33a)}\\
\text{subj. to} \quad & \mathbf{u} \in \mathcal{U}^{N_\delta} && \text{(33b)}\\
& Gx_0 + H\mathbf{u} \in \mathcal{K}^{N_\delta+1} && \text{(33c)}\\
& r_{d_i}^\top x_0 = r_{d_i}^\top y_{ij}. && \text{(33d)}
\end{aligned}$$

Notice that we no longer fix x_0; rather, we implicitly look for a point on the set $\{x_0 \mid r_{d_i}^\top x_0 = r_{d_i}^\top y_{ij}\} \cap \mathcal{K}$, when y_{ij} varies, that is a feasible point. We also do not erode the constraints as we did before since our goal here is find an *over-approximation* of the true kernel.

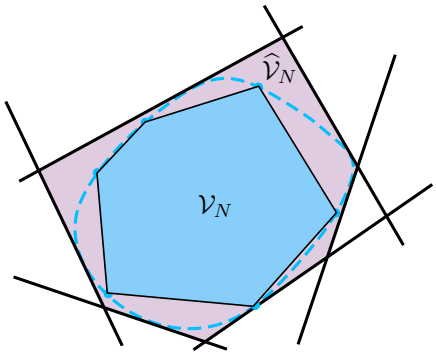

Figure 7: The viability kernel is sandwiched in between the under-approximation $\mathcal{V}_N^{\zeta,\epsilon}$ and the over-approximation $\widehat{\mathcal{V}}_N^{\zeta,\epsilon_{\mathrm{o}}}$.

Once a feasible solution to (33) is found for the desired accuracy ϵ_{o} of the bisection search, we stop the iterations and store two entities:

(S1) The last *infeasible* step, i.e. the last value of y_{ij} for which (33) was infeasible. Denote this value by $y_{ij}^{\inf *}$;

(S2) The feasible solution pair $(\mathbf{u}_i^*, x_{0i}^*)$ (indexed by i to correspond to the direction r_{d_i}).

We first use entity (S1) to form our over-approximation along r_{d_i} as the halfspace

$$\{x \mid r_{d_i}^\top x \leq r_{d_i}^\top y_{ij}^{\inf *}\}. \tag{34}$$

Clearly the set $\{x \mid r_{d_i}^\top x = r_{d_i}^\top y_{ij}^{\inf *}\}$ is a supporting hyperplane of $\mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$ with an arbitrary (and desirably conservative) error ϵ_{o}. By repeating the above procedure for all N directions and forming the set

$$\widehat{\mathcal{V}}_N^{\zeta,\epsilon_{\mathrm{o}}} := \bigcap_{i=1}^{N} \{x \mid r_{d_i}^\top x \leq r_{d_i}^\top y_{ij}^{\inf *}\} \tag{35}$$

we obtain (Fig. 7)

$$\mathcal{V}_N^{\zeta,\epsilon} \subseteq \mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K}) \subseteq \widehat{\mathcal{V}}_N^{\zeta,\epsilon_{\mathrm{o}}}. \tag{36}$$

The error $\mathrm{vol}(\mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})) - \mathrm{vol}(\mathcal{V}_N^{\zeta,\epsilon})$ of our main under-approximation algorithm can then be quantitatively bounded above as $\mathrm{vol}(\widehat{\mathcal{V}}_N^{\zeta,\epsilon_{\mathrm{o}}}) - \mathrm{vol}(\mathcal{V}_N^{\zeta,\epsilon})$. This upper-bound monotonically decreases as N increases, and converges to a numerical constant as $\epsilon, \epsilon_{\mathrm{o}} \to 0$ and $N, \zeta \to \infty$.

The entity (S2) from the feasibility program (33) can help us improve our under-approximation. This is discussed next.

7 Improving the Under-Approximation

We describe two techniques that help improve the quality of our under-approximation for finite number of samples.

7.1 Center of Mass & Gauss-Kronecker Curvature

If we were to naively sample from a uniform distribution on $\mathcal{S}_2^{n-1}$ to compute our under-approximating set, two elements would negatively affect the quality of such approximation: (a) the unbalanced distances between the unit ball centered at v_0 and the boundary points of the true kernel, i.e. the distance between v_0 and the center of mass of the kernel, and (b) the regions of the boundary of the kernel with high curvature. This is because a uniform distribution

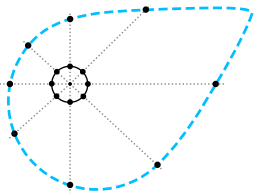

Figure 8: Mapping a uniform distribution on $\mathcal{S}_2^{n-1} \subseteq \mathcal{C}$ linearly onto $\partial\mathcal{C}$ may result in a distribution that is non-uniform.

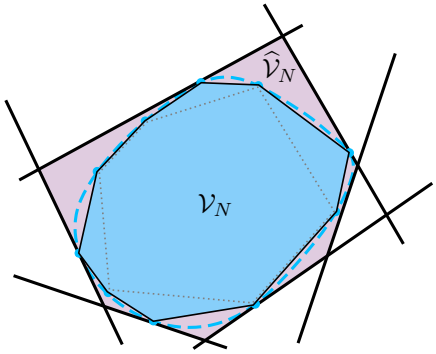

Figure 9: Injecting additional under-approximation steps guided by the available over-approximation facets improves the quality of the resulting set (compare to Fig. 7). Areas of the kernel with high curvature are now covered at a faster rate.

on $\mathcal{S}_2^{n-1}$ mapped onto $\partial\,\mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$ could yield a distribution that is far from uniform depending on severity of (a) or (b); Fig. 8.

The former issue can be somewhat mitigated by continually moving the point v_0 to the center (e.g. in the sense of Chebyshev) of the set $\mathcal{V}_N^{\zeta,\epsilon}$ every time a new vertex is added. The latter issue can be addressed using the information obtained from the over-approximation procedure discussed in the previous section, specifically using the stored entity (S2). The idea is that the greater the distance between the ith under-approximation vertex and over-approximation facet, the higher the Gauss curvature of the boundary of the true kernel in the neighborhood of that unexplored region.

To account for this problem, we perform an additional step after the ith iteration of our combined algorithm, every time a vertex is added to the under-approximation and an over-approximating halfspace is formed along the direction r_{d_i}: The stored value x_{0i}^* in (S2) approximates the *support vector* (the point in which the support function of a convex set touches its boundary) of the viability kernel in the direction r_{d_i} with ϵ_{o} accuracy. Therefore, we use x_{0i}^* and execute a single instance of our under-approximation procedure this time along not r_{d_i}, but along the direction $x_{0i}^* - v_0$ (corresponding to the ray passing through v_0 and x_{0i}^*).

By doing so, we allow the over-approximation to "guide" where we should look for the next under-approximation vertex. Moreover, for this new instance of the under-approximation algorithm we can limit the bisection search to a diameter of roughly $2\epsilon_{\mathrm{o}}$ around x_{0i}^* (instead of perfoming the search between v_0 and the point at which the new ray intersects $\partial\mathcal{K}$) since we know that x_{0i}^* is already fairly close to the boundary of the true kernel. The resulting vertex is not only in close proximity to the point where the over-approximating halfplane at the ith iteration has been formed (thus providing a superior confidence that the viability kernel is sandwiched tightly in that area), but also covers the parts of the boundary that could potentially have a high Gauss curvature; Fig. 9.

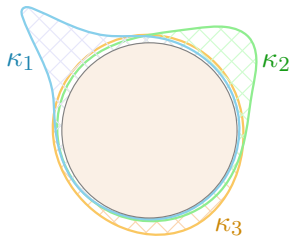

Figure 10: The vMF density function on the unit ball for three different mean directions and concentrations with $\kappa_1 > \kappa_2 > \kappa_3$.

7.2 Biased Random Sampling

The shortcomings of uniform sampling are more pronounced in high dimensions. Thus we additionally seek to *bias* the distribution on the unit ball to mitigate these shortcomings. To this end, we present a few heuristic techniques that are still based on random sampling so as to keep the optimality results of Section 4.2, but could potentially improve the performance of the algorithm.

We will make use of the von-Mises Fisher (vMF) distribution [36] whose density function is given by

$$f_{\text{vMF}}(x; \mu, \kappa) := C(\kappa) e^{\kappa \mu^\top x} \tag{37}$$

with *concentration* $\kappa \geq 0$, mean direction μ ($\|\mu\| = 1$), and a normalizing constant $C(\kappa)$. The parameter κ determines how samples drawn from this distribution are concentrated around the mean direction μ. For $\kappa = 0$ the vMF reduces to the uniform density; otherwise, it resembles a normal density (with compact support on the unit ball), centered at μ with variance inversely proportional to κ (Fig. 10). As $\kappa \to \infty$ the vMF converges to a point distribution.

The solutions we propose are by no means exhaustive, and there may be better ways of guiding the sampling process depending on the problem in hand or if we have some *a priori* knowledge of the shape of the viability kernel.

7.2.1 Gradient-Like Methods

The first approach we discuss is related to how some measure of the error varies in consecutive steps of our combined under- and over-approximation algorithm. We let this change dictate from what vMF distribution should the next sample be drawn.

Let $\mathcal{V}_i^{\zeta,\epsilon}$ and $\widehat{\mathcal{V}}_i^{\zeta,\epsilon_o}$ respectively denote the under-approximation and over-approximation of the viability kernel after adding the ith sample direction r_{d_i}, and let the state space be equipped with some metric d. Define the error function (possibly nonconvex) $\text{err}\colon \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ at this ith iteration as

$$\text{err}(r_{d_i}) := d(\mathcal{V}_i^{\zeta,\epsilon}, \widehat{\mathcal{V}}_i^{\zeta,\epsilon_o}). \tag{38}$$

Ideally, the metric d is chosen such that $\text{err}(r_{d_i})$ is monotonically non-increasing as i increases (i.e. as we produce more accurate approximations). Define the normalized quantity

$$\nabla \text{err}(r_{d_i}) := \frac{1 - d(\mathcal{V}_i^{\zeta,\epsilon}, \widehat{\mathcal{V}}_i^{\zeta,\epsilon_o})/d(\mathcal{V}_{i-1}^{\zeta,\epsilon}, \widehat{\mathcal{V}}_{i-1}^{\zeta,\epsilon_o})}{\frac{1}{\pi}\cos^{-1}\left(\langle r_{d_i}, r_{d_{i-1}} \rangle\right)}, \tag{39}$$

treating the pathological case $\frac{0}{0}$ as 0. The magnitude of $\nabla \text{err}(r_{d_i})$ approximates the rate of improvement between two consecutive iterations of our combined algorithm. Thus, we can use this information to draw the next sampling direction $r_{d_{i+1}}$ as

$$r_{d_{i+1}} \sim f_{\text{vMF}}(x; \mu_{i+1}, \kappa_{i+1}) \tag{40}$$

with mean and concentration parameters updated as

$$\kappa_{i+1} = \max\{\nu_0, \omega_i\}, \tag{41}$$

$$\mu_{i+1} = \text{sgn}(-\nu_0 + \omega_i) \cdot r_{d_i}, \tag{42}$$

where

$$\omega_i := \nu_1 \tanh\left(\nu_2 \nabla \operatorname{err}(r_{d_i})\right). \tag{43}$$

Here the positive scalars ν_0, ν_1, and ν_2 are design choices; for example, ν_2 dictates how quickly the tangent hyperbolic function reaches its maximum ν_1 where it levels off, while ν_0 (which is a small scalar) limits how close the resulting vMF distribution can be to the uniform one and, when necessary, sets the sign of the direction μ_{i+1} of the distribution to the negative of r_{d_i}.

The intuition behind such an update rule is that when sampling along a given direction r_{d_i} causes the magnitude of $\nabla \operatorname{err}$ to become large (thus causing the error to decrease quickly), drawing the next sample from a distribution that is concentrated (proportional to this change) around the same direction may be a good choice to continue to reduce the error: $\kappa_{i+1} = \omega_i$ and $\mu_{i+1} = r_{d_i}$. If, on the other hand, the magnitude of $\nabla \operatorname{err}$ is small, we are approaching a local minima of the quasilinear function $\operatorname{err}(\cdot)$. Therefore, we would want to sample from a distribution that is *not* concentrated around the current direction r_{d_i}, and could also be such that it potentially minimizes the likelihood of sampling around r_{d_i}. In such a case, we would bias the sampling to be slightly concentrated around the opposite direction (and also decrease the likelihood of resampling around r_{d_i}) so as to attempt to get us out of this local minima, while not making it to be too far off from a uniform distribution: $\kappa_{i+1} = \nu_0$ and $\mu_{i+1} = -r_{d_i}$.

Volumetric Measure: We can use volume as a metric in (38) so that $\operatorname{err}(r_{d_i}) := \operatorname{vol}(\widehat{\mathcal{V}}_i^{\zeta,\epsilon_\mathrm{o}}) - \operatorname{vol}(\mathcal{V}_i^{\zeta,\epsilon})$. Of course, computing the volume of a polytope is #P-hard [37]. However, we can conservatively approximate $\operatorname{err}(r_{d_i})$ by calculating analytically the volume of appropriate inscribed and circumscribed ellipsoids: A minimum volume circumscribed ellipsoid (mVCE) containing a polytope that is represented by its vertices can be computed efficiently via a semidefinite program [38]. The same is true for a maximum volume inscribe ellipsoid (MVIE) that is contained in a polytope represented by its facets. (Note that the converse problems are NP-hard.) If we computed the mVCE of $\mathcal{V}_i^{\zeta,\epsilon}$ and shrunk it by a factor of n (the dimension), then the resulting ellipsoid would be a subset of $\mathcal{V}_i^{\zeta,\epsilon}$. Similarly, if we enlarged the MVIE of $\widehat{\mathcal{V}}_i^{\zeta,\epsilon_\mathrm{o}}$ by a factor of n, then $\widehat{\mathcal{V}}_i^{\zeta,\epsilon_\mathrm{o}}$ would be a subset of the resulting ellipsoid. The difference in the volume of these two ellipsoids would be an upper-bound on $\operatorname{err}(r_{d_i})$.

Working with the volume of these extremal ellipsoids does preserve order, in that the relations $\operatorname{vol}(\mathcal{V}_{i+1}^{\zeta,\epsilon}) \geq \operatorname{vol}(\mathcal{V}_i^{\zeta,\epsilon})$ and $\operatorname{vol}(\widehat{\mathcal{V}}_{i+1}^{\zeta,\epsilon_\mathrm{o}}) \leq \operatorname{vol}(\widehat{\mathcal{V}}_i^{\zeta,\epsilon_\mathrm{o}})$ are also true for their extremal ellipsoids. The reason is that adding constraints (additional vertices or facets) to a convex optimization problem (the SDPs) cannot decrease the value of the optimal objective functions (correlated with ellipsoid volumes). As a result, the error upper-bound, just like the error itself, is monotonically non-increasing as i increases. A disadvantage of working with scaled extremal ellipsoids is that the shrinkage/enlargement by a factor of n can be too conservative particularly for larger n.

Hausdorff Distance: The Hausdorff distance between two compact convex sets $\mathcal{C}_1$ and $\mathcal{C}_2$ in $\mathbb{R}^n$ in terms of their support functions is defined as

$$d_\mathrm{H}(\mathcal{C}_1, \mathcal{C}_2) := \max_{\ell \in \mathcal{S}^{n-1}} \left\{ |\rho_{\mathcal{C}_1}(\ell) - \rho_{\mathcal{C}_2}(\ell)| \right\}. \tag{44}$$

We can use this metric to define the error: $\operatorname{err}(r_{d_i}) = d_\mathrm{H}(\mathcal{V}_i^{\zeta,\epsilon}, \widehat{\mathcal{V}}_i^{\zeta,\epsilon_\mathrm{o}})$. Computing this distance between $\mathcal{V}_i^{\zeta,\epsilon} = \operatorname{conv}(\{v_j\}_{j=0}^i)$ and $\widehat{\mathcal{V}}_i^{\zeta,\epsilon_\mathrm{o}} =: P_i x \leq p_i$ (where P and p respectively are the appropriate matrix and vector corresponding to the facet-based outer polytope defined in (35) but with i faces) can be cast as a series of LPs: For a fixed direction $\ell \in \mathcal{S}^{n-1}$, we compute at every iteration i

$$\rho_{\mathcal{V}_i^{\zeta,\epsilon}}(\ell) = \max_{x,\, \{\lambda_j \geq 0\}} \quad \ell^\top x \tag{45a}$$

$$\text{subj. to} \quad x = \sum_{j=0}^{i} \lambda_j v_j, \;\; \sum_{j=0}^{i} \lambda_j = 1, \tag{45b}$$

as well as

$$\rho_{\widehat{\mathcal{V}}_i^{\zeta,\epsilon_\mathrm{o}}}(\ell) = \max_{x} \quad \ell^\top x \tag{46a}$$

$$\text{subj. to} \quad P_i x \leq p_i. \tag{46b}$$

The Hausdorff distance can then be approximated over a finite number of directions $\mathcal{L} \subset \mathcal{S}^{n-1}$ as

$$\hat{d}_{\mathrm{H}}(\mathcal{V}_i^{\zeta,\epsilon}, \widehat{\mathcal{V}}_i^{\zeta,\epsilon_{\mathrm{o}}}) := \max_{\ell \in \mathcal{L}} \left\{ \left| \rho_{\widehat{\mathcal{V}}_i^{\zeta,\epsilon_{\mathrm{o}}}}(\ell) - \rho_{\mathcal{V}_i^{\zeta,\epsilon}}(\ell) \right| \right\}$$
$$\leq \max_{\ell \in \mathcal{S}^{n-1}} \left\{ \left| \rho_{\widehat{\mathcal{V}}_i^{\zeta,\epsilon_{\mathrm{o}}}}(\ell) - \rho_{\mathcal{V}_i^{\zeta,\epsilon}}(\ell) \right| \right\} = d_{\mathrm{H}}(\mathcal{V}_i^{\zeta,\epsilon}, \widehat{\mathcal{V}}_i^{\zeta,\epsilon_{\mathrm{o}}}). \quad (47)$$

This approximation is not conservative due to the inequality in (47). But as far as guiding the sampling, it may yield a viable alternative to the volumetric error described above (that could be excessively conservative). Unfortunately, performing these additional LPs at the end of each iteration of our algorithm may undermine its efficiency.

7.2.2 Purely Heuristic Methods

We also propose two purely heuristic methods as alternatives to the gradient-like methods presented above.

Averaged Opposite Direction: For this approach we associate a given vMF density function to the negative of *each* individual direction vector we have generated so far, and draw our next sampling direction from a convolution of these density functions. That is, in some sense we are drawing at random a sample whose expected value lies in the opposite direction of the samples we have already drawn. The intuition here is that we can approach "true" uniformity at a higher rate by sampling in the direction whose neighborhood we have not yet sampled as densely.

We can vary the concentration of the vMF densities as an increasing function of the iteration step i (e.g. linearly with the number of vertices), such that at the beginning steps of our algorithm these densities are closer to uniform, and as we progress and generate more vertices/directions their concentrations increase.[7] A design parameter ν_1 can be a multiplier for the number of vertices. We cap the concentration of the distributions to 100 so as to not lose the randomness properties of the algorithm when i grows too large.

Point-to-Plane Distance: This approach is based on simply calculating the smallest distance from every vertex of $\mathcal{V}_i^{\zeta,\epsilon}$ to facets of $\widehat{\mathcal{V}}_i^{\zeta,\epsilon_{\mathrm{o}}}$, and identifying the largest of such distances. This quantity in some sense describes the gaps between the two sets over which we are most uncertain about the boundary of the true kernel.

Let $\bar{v}_j$ be the vertex in $\mathcal{V}_i^{\zeta,\epsilon} = \mathrm{conv}(\{v_j\}_{j=0}^i)$ with the largest point-to-plane distance. To guide the sampling at the next iteration of the algorithm, we calculate the ray that passes through the center v_0 and the point in the facet of $\widehat{\mathcal{V}}_i^{\zeta,\epsilon_{\mathrm{o}}}$ that is the closest to $\bar{v}_j$. To determine our next sampling direction $r_{d_{i+1}}$, we then sample from a vMF distribution whose mean is the unit vector along this ray and whose concentration is dependant on the calculated point-to-plane distance from $\bar{v}_j$ to $\widehat{\mathcal{V}}_i^{\zeta,\epsilon_{\mathrm{o}}}$. A design parameter can again be a multiplier ν_1 for this distance.

7.2.3 Calibration and Comparison on Random Systems

To compare the performance of the four guiding methods described above, we first roughly calibrated the design parameters for each approach by running each guiding technique on 30 randomly generated systems across 2D, 3D, and 4D state and input dimensions[8] for a variety of exponents of 10. For example, we tested the volume extremal ellipsoids and the Hausdorff distance methods for all combinations of $\nu_0 = 10^{\{-1,0\}}$, $\nu_1 = 10^{\{0,1,2\}}$, and $\nu_2 = 10^{\{0,1\}}/n$ (made dimension n dependent so that the tangent hyperbolic reaches its maximum ν_1 more slowly when n is larger); and the averaged opposite direction and the point-to-plane distance methods for $\nu_1 = 10^{\{-2,-1,0,1\}}/n$ (made dimension dependent so that the vMF concentration increases more slowly when n is larger). We found the optimal parameters to be $\nu_0 = 0.1$, $\nu_1 = 1$, $\nu_2 = 1/n$ for the volume extremal ellipsoids; $\nu_0 = 1$, $\nu_1 = 1$, $\nu_2 = 1/n$ for the Hausdorff distance; $\nu_1 = 0.1/n$ for the averaged opposite direction; and $\nu_1 = 1/n$ for the point-to-plane distance. We emphasize that this calibration was meant only to achieve parameters that were roughly the correct order of magnitude; we did no further fine tuning.

With the calibrated methods in hand, we then ran all four methods on a fresh set of 60 randomly generated systems across 3D, 4D, and 5D state and input dimensions. We were then able to evaluate the performance of each method by taking advantage of the fact that all the approximations generated are under-approximations. Thus, the under-approximation with the largest possible volume must be (in a volumetric sense) the closest to the true viability kernel.

[7] The density function of the convolution is closer to uniform if (a) the individual densities are closer to uniform, or (b) the existing directions all cancel each other out.

[8] Our tests were limited by 5D due to the need to directly compute the volume of polytopes for performance assessment.

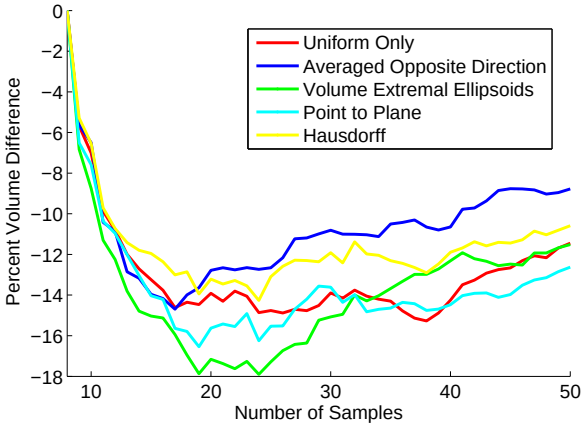

Figure 11: Percent difference in volume between the result of each guiding method and the best result out of all the guiding methods for a given number of vertices (averaged over 60 randomly generated systems across 3, 4, and 5 state and input dimensions).

Using this fact we found for each random system and for a given number of vertices the best under-approximation selected from the results of all four different methods (as well as the basic uniform sampling method). Then, for each system and number of vertices we calculated the percent difference in volume between the best under-approximation and the result of each method. (Note that this percent difference must be negative, since the result from each method must have a volume smaller than or equal to the result of the best method.) Thus a method with a less negative value, i.e. closer to 0, indicates a method that performs better, in a volumetric sense, than a method with a more negative value (which indicates a much smaller and thus less accurate under-approximation in a volumetric sense). Fig. 11 plots the resulting percent volume difference by method as a function of the number of vertices, averaged over the 60 randomly generated systems.

On average, the averaged opposite direction (blue) outperforms all methods. More specifically, it improves the resulting under-approximation volume by about 5% over the uniform sampling (red). The Hausdorff distance (yellow) performs consistently better than uniform sampling but since we only compute a crude approximation of the actual Hausdorff distance in (47) (at every step we only use the directions along which we have already generated a vertex so as to maintain monotonicity of the distance approximate while keeping the computation times on par with the other methods), this improvement is expected to be more emphasized for higher number of samples than only 50. The volume extremal ellipsoids (green) initially performs far better than uniform sampling in 3D and 4D, but then quickly degrades in 5D due to shrinkage/expansion of the ellipsoids by a factor of n that is much more exaggerated in higher dimensions. The point-to-plane distance (cyan) simply does not show any improvement over uniform sampling, at least on average in our test setup. We will use the averaged opposite direction as our guiding method of choice for our examples in Section 8. However, we do emphasize again that the performance of these heuristics is problem-dependent. Therefore, while in our limited random tests one method might have outperformed others on average, it may be that another method is even more suitable and performs significantly better for a given problem.

8 Examples

8.1 The Double-integrator

First, consider the simple dynamics $\ddot{x} = u$ with $\delta = 0.05\,\text{s}$. The constraints $\mathcal{K} = \{x \mid \|x\|_\infty \leq 0.5\}$ and $\mathcal{U} = [-0.15, 0.15]$ are to be respected over $\mathbb{T} = [0, 1]$. To find a bound M on the vector field and construct the eroded set $\mathcal{K}_\downarrow(M, \delta)$ we use the following procedure: We first scale the state space by computing for every dimension d the quantity $z^*(\mathbf{e}_d) - z^*(-\mathbf{e}_d)$ with $z^*(\mathbf{e}_d) := \arg\max_z \mathbf{e}_d^\top z$ subject to $z = \dot{x}$ and $x \in \mathcal{K}$, $u \in \mathcal{U}$, and where $\mathbf{e}_d$ denotes

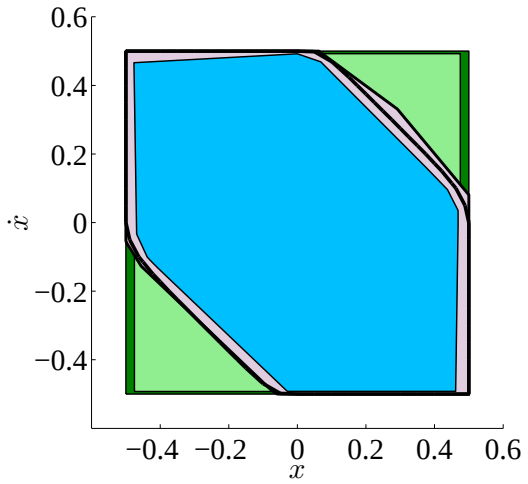

Figure 12: Polytopic under-approximation $\mathcal{V}_N^{\zeta,\epsilon}$ (blue) and over-approximation $\widehat{\mathcal{V}}_{N/2}^{\zeta,\epsilon_o}$ (lavender) of $\mathrm{Viab}_{\mathbb{T}}^{\mathrm{sd}}(\mathcal{K})$ for the double-integrator example with $N = 20$ samples. The sets $\mathcal{K}$ and $\mathcal{K}_{\downarrow}(M,\delta)$ are shown in dark and light green, respectively. The SD level-set approximation [7] is also shown (outlined in thick black line).

the standard basis vector spanning dth dimension. Then we divide all such quantities by their minimum value among all dimensions, and transform the dynamics accordingly so that a state increment in all dimensions is equivalent. At this stage, we can calculate M as the optimal value of $\max_{x,u} \|\dot{x}\|_\infty$ subject to $x \in \mathcal{K}$ and $u \in \mathcal{U}$. The eroded set $\mathcal{K}_{\downarrow}(M,\delta)$ is now constructed in the scaled state space according to Lemma 1.

We used $\zeta = 4$th order discretization and $\epsilon = \epsilon_o = 0.01$-accurate bisection search to obtain the under- and over-approximations shown in Fig. 12. The under-approximation is computed using $N = 20$ randomly generated samples via Algorithm 1. The sampling process was guided via the techniques in Section 7. The over-approximation polytope consists of $N/2$ facets and is computed according to Section 6. The overall computation time was 18 s.

8.2 12D Quadrotor Flight Envelope Protection

We now evaluate our algorithm on the benchmark example described in [39]. Consider the full-order model of a quadrotor based on the nonlinear Newton-Euler rigid body equations of motion. The state vector

$$x = \begin{bmatrix} \mathrm{x} & \mathrm{y} & \mathrm{z} & \dot{\mathrm{x}} & \dot{\mathrm{y}} & \dot{\mathrm{z}} & \phi & \theta & \psi & \dot{\phi} & \dot{\theta} & \dot{\psi} \end{bmatrix}^\top \in \mathbb{R}^{12} \tag{48}$$

is comprised of translational positions in [m] with respect to a global origin, their derivatives (linear velocities in x, y, z directions) in [m/s], the Eulerian angles roll ϕ, pitch θ, and yaw ψ in [rad], and their respective derivatives (angular velocities) in [rad/s]. The control input is the vector $u = [u_1\ u_2\ u_3\ u_4]^\top \in \mathbb{R}^4$ consisting, respectively, of the total thrust in [m/s^2] normalized with respect to the mass of the quadrotor (u_1) and the second-order derivatives $\ddot{\phi}$, $\ddot{\theta}$, $\ddot{\psi}$ of the Eulerian angles in [rad/s^2] (u_2 through u_4). The system is under-actuated since there are six degrees of freedom but only four actuators. By linearizing the equations of motion about the hover condition $\phi = 0$, $\theta = 0$, and $u_1 = g$ (with $g \approx 9.81$ being the acceleration of gravity) one would obtain the model $\dot{x} = Ax + Bu$ with the state and the input now representing *deviation* from the equilibrium. We follow the example detailed by [40] of an agile quadrotor in which the state is sampled at a frequency of 10 Hz. (See the same reference for values of the system matrices A and B.)

For safe operation of the vehicle the Eulerian angles ϕ and θ and the speed profile $V := \|[\dot{\mathrm{x}}\ \dot{\mathrm{y}}\ \dot{\mathrm{z}}]\|$ are bounded as $\phi, \theta \in [-\frac{\pi}{4}, \frac{\pi}{4}]$ and $V \leq 5$. The angular velocities are constrained as $\dot{\phi}, \dot{\theta}, \dot{\psi} \in [-3, 3]$. We further assume that the vehicle must safely fly within the range of 1 to 7 m above the ground in z direction in an environment that stretches 6 m in each direction in the x-y plane. These constraints form the flight envelope $\mathcal{K}$. The vector u is constrained by the hyper-rectangle $\mathcal{U} := [-g, 2.38] \times [-0.5, 0.5]^3$.

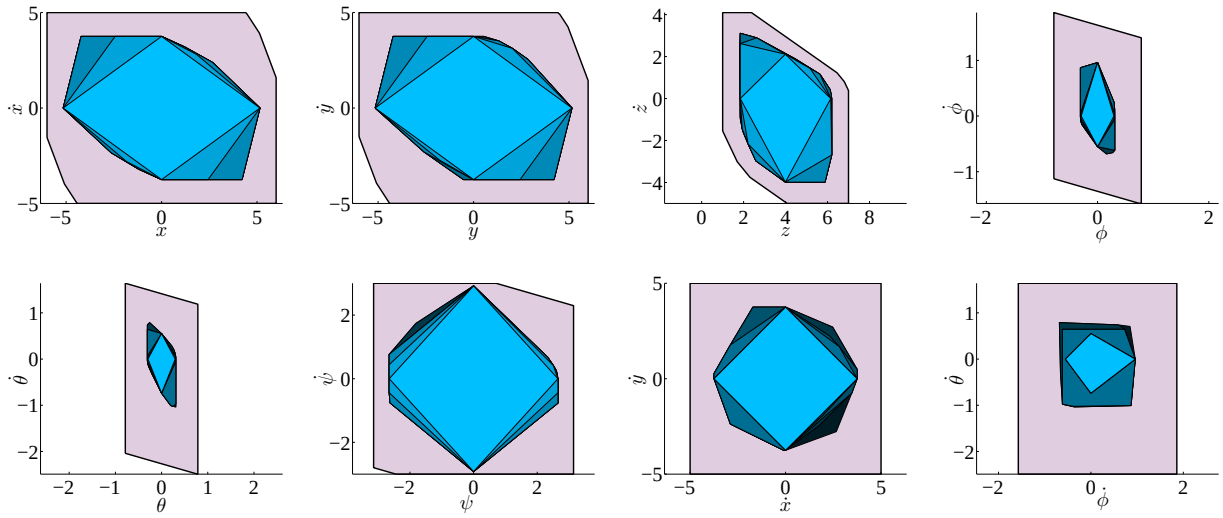

Figure 13: Selected 2D projections of the polytopic approximation of $\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}(\mathcal{K})$ for the flight envelope example. Under-approximations for $N = 24, 48, 96, 500, 1000, 2000, 3000$ vertices are shown with $N = 24$ in the lightest shade of blue (innermost set), and $N = 3000$ in the darkest shade of blue. An over-approximation (outermost set) with $N/2$ facets is also shown in lavender. The approximations are tight and touch each other to within a constant accuracy in at least $N/2$ points—a fact that is obscured by the projections.

The quadrotor can travel a distance of roughly half a meter between any two consecutive sampling times despite the relatively high sampling frequency. This fact further warrants the treatment of such a safety-critical system through a SD framework. We wish to compute the set of initial states for which safety can be maintained over $\mathbb{T} = [0, 2]$.

We warm-start our approximation algorithm (in the scaled state space) by first sampling along all axes in order to obtain a full-dimensional object, and then along 72 uniformly-spaced vectors in x_i-x_{i+3} and x_{i+6}-x_{i+9} subspaces for $i = 1, 2, 3$. The remaining samples are generated randomly by guiding the vMF sampler via the averaged opposite direction method. The discretization order and the bisection search accuracies are the same as in the previous example.

Fig. 13 shows selected 2D projections of our sampling-based polytopic approximations of $\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}(\mathcal{K})$ in the original unscaled state space. The algorithm requires about 3 s (without optimizing the code for speed) to generate a new vertex of the under-approximation and 5 s to generate a facet of the over-approximation. The under-approximation is tight in the sense that each vertex of the polytope belongs to the boundary of the true viability kernel with some *a priori* known accuracy due to the SD nature of the problem. The over-approximation is also tight in that each facet of the polytope touches the boundary of the true kernel in at least one point. The two approximating sets sandwich the boundary of the viability kernel to within a certain precision in at least half of the number of vertices of the under-approximation, providing an added layer of confidence about the precise location of the kernel. The tightness of the sets are unfortunately unobservable in the projection plots.

9 Conclusions, Extensions, and Future Work

We presented a scalable sampling-based algorithm to generate tight under- and over-approximations of the viability kernel under SD LTI dynamics. We provided correctness and convergence results, discussed a number of heuristics to bias the sampling process for improved performance, and demonstrated the algorithm on a 12D problem of flight envelope protection for an autonomous quadrotor.

The extension to discrete-time systems is straightforward: The set $\mathcal{K}$ is directly used in the feasibility program (24) and thus the vertices of the resulting under-approximation and the facets of the over-approximation touch the boundary of the true viability kernel exactly. The algorithm can also be extended to piecewise affine SD (or discrete-time) system

by recasting the optimization problem as a mixed-integer program. Due to space limitations we will present our results on the synthesis of the safety-preserving controllers as well as sufficient conditions under which a generated under-approximation is controlled-invariant (and thus could be used to enforce infinite-horizon safety) in a separate paper.

Robustifying our analysis against unknown but bounded disturbances is challenging (due to the minimax nature of the problem), but a sufficient solution is straightforward and can be done in a number of ways. For instance, the disturbance set could be propagated forward in time, according to which the constraint set could be further eroded to take into account the effect of this uncertainty in future time steps. Alternatively, a pre-stabilization technique [41] could be used, under certain additional assumptions, such that a portion of the control is dedicated to managing the growth of the disturbance set propagation while the remaining portion is employed to keep the trajectory in $\mathcal{K}$.

The simulation based nature of our algorithm readily admits incorporation of time delays. We will make our initial results on SD systems with transport delays available as a technical report for the interested reader. Finally, we also mention that a similar sampling-based technique can be formulated to approximate nonconvex maximal reachable tubes under LTI dynamics, leveraging the fact that the underlying reachable sets (time slices of the tube) are convex.

Future work for piecewise affine systems includes finding alternative ways to solve the mixed-integer program so as to improve efficiency. The extension of the algorithm to nonconvex constraints remains another challenging problem.

A Appendix: Proof of Proposition 1

To prove Proposition 1 we will introduce and make use of a few well-known results from the literature of random algorithms for estimating the boundary of a convex body.

Lemma 3 ([33])**.** *Let $\mathcal{C}$ be a convex body in $\mathbb{R}^n$ with $\partial\mathcal{C}$ C^2, let $f : \partial\mathcal{C} \to \mathbb{R}_+$ be a probability density function defined on $\partial\mathcal{C}$, let $\mathbb{P}_f$ be the probability measure defined by f, and let $\mathbb{E}(f, N)$ be the expected volume of the convex hull of N points chosen randomly on $\partial\mathcal{C}$ with respect to $\mathbb{P}_f$. Then*

$$\lim_{N\to\infty} \left(\mathrm{vol}(\mathcal{C}) - \mathbb{E}(f, N)\right) N^{\frac{2}{n-1}} = c_n(\mathcal{C}), \tag{49}$$

where $c_n(\mathcal{C})$ is a constant which depends only on the dimension n, the distribution of f, and the shape of $\mathcal{C}$.

To use this result, we also need the following lemma, which introduces a fictitious source of error between the outcome of our algorithm and the viability kernel:

Lemma 4 ([42])**.** *For every compact convex set $\mathcal{C}$, there exists a compact convex set $\mathcal{C}' \subseteq \mathcal{C}$ whose boundary $\partial\mathcal{C}'$ is C^2, and $\mathrm{vol}(\mathcal{C}) - \mathrm{vol}(\mathcal{C}') = \epsilon_{\mathrm{smooth}}$ for some arbitrarily small positive scalar $\epsilon_{\mathrm{smooth}}$.*

Define via Lemma 4, a convex body $\overline{\mathrm{Viab}}'$ that is a C^2 approximation of $\overline{\mathrm{Viab}}^{\mathrm{sd}}_{\mathbb{T}}(\mathcal{K}, \zeta, \epsilon)$ such that

$$\mathrm{vol}(\overline{\mathrm{Viab}}^{\mathrm{sd}}_{\mathbb{T}}(\mathcal{K}, \zeta, \epsilon)) - \mathrm{vol}(\overline{\mathrm{Viab}}') = \epsilon_{\mathrm{smooth}}. \tag{50}$$

Lemma 5. *Any compact convex set $\mathcal{C}$ with $r_0 \in \mathcal{C}$ is homeomorphic to $\mathcal{B}^n_2(r_0, 1)$ [43], and in particular we can define the invertible mapping $m : S^{n-1}_2(r_0) \to \partial\mathcal{C}$ as $m(r_d) \equiv$* BISECTION-FEASIBILITY$\big(r_0,$ FIND-INTERSECTION-ON-BOUNDARY$(\mathcal{C}, \vec{r}), \mathcal{C}\big)$*, where $\vec{r}$ has origin $r_0 \in \mathcal{C}$ and direction r_d.*

We are now ready to prove Proposition 1.

Proof of Proposition 1. Let $f(x)$ be the probability distribution used by SAMPLE-RAY(v_0) to generate samples on the unit sphere. Then by using the mapping m from Lemma 5, we can perform a change of variables to define a new probability density function $g(m(r_d))$ on $\partial\overline{\mathrm{Viab}}'$ [44]. Then by Lemma 3 we have

$$\lim_{N\to\infty} \Big(\mathrm{vol}(\overline{\mathrm{Viab}}') - \mathbb{E}(g, N)\Big) N^{\frac{2}{n-1}} = \tilde{c}_n(\overline{\mathrm{Viab}}') \tag{51}$$

for some constant $\tilde{c}_n(\overline{\mathrm{Viab}}')$ which depends only on the dimension n, the distribution g, and the shape of $\overline{\mathrm{Viab}}'$.

Since by Lemma 4 $\overline{\mathrm{Viab}}'$ can be made arbitrarily close to $\overline{\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}}(\mathcal{K},\zeta,\epsilon)$ (i.e. we can take $\epsilon_{\mathrm{smooth}} \to 0$), the distribution g is mapped almost identically on $\partial\overline{\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}}(\mathcal{K},\zeta,\epsilon)$ and we can write (51) as

$$\begin{aligned}\lim_{N\to\infty}\left(\mathrm{vol}(\overline{\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}}(\mathcal{K},\zeta,\epsilon)) - \mathbb{E}(g,N)\right) N^{\frac{2}{n-1}} \\ = \lim_{N\to\infty}\left(\mathrm{vol}(\overline{\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}}(\mathcal{K},\zeta,\epsilon)) - \mathrm{vol}(\mathcal{V}_N^{\zeta,\epsilon})\right) N^{\frac{2}{n-1}} \\ = \tilde{c}_n(\overline{\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}}(\mathcal{K},\zeta,\epsilon)) \quad (52)\end{aligned}$$

since $\mathbb{E}(g,N) = \mathrm{vol}(\mathcal{V}_N^{\zeta,\epsilon})$.

Taking the limit as $\zeta \to \infty$ and $\epsilon \to 0$ on (52) gives

$$\lim_{\substack{N\to\infty\\ \zeta\to\infty\\ \epsilon\to 0}}\left(\mathrm{vol}(\overline{\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}}(\mathcal{K})) - \mathrm{vol}(\mathcal{V}_N^{\zeta,\epsilon})\right) N^{\frac{2}{n-1}} = \tilde{c}_n(\overline{\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}}(\mathcal{K})). \quad (53)$$

Now, by (27), we can replace $\mathrm{vol}(\overline{\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}}(\mathcal{K}))$ to get

$$\lim_{\substack{N\to\infty\\ \zeta\to\infty\\ \epsilon\to 0}}\left(\mathrm{vol}(\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}(\mathcal{K})) - \epsilon_{\mathrm{cont}}(M\delta) - \mathrm{vol}(\mathcal{V}_N^{\zeta,\epsilon})\right) N^{\frac{2}{n-1}} = \tilde{c}_n(\overline{\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}}(\mathcal{K})). \quad (54)$$

The shape of $\overline{\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}}(\mathcal{K})$ depends only on the shape of $\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}(\mathcal{K})$ and the value $M\delta$. Thus there exists an appropriate function c_n such that $\tilde{c}_n(\overline{\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}}(\mathcal{K})) = c_n(\mathrm{Viab}^{\mathrm{sd}}_{\mathbb{T}}(\mathcal{K}), M\delta)$. Consequently, we arrive at (30). □

References

[1] J. H. Gillula, S. Kaynama, and C. J. Tomlin, "Sampling-based approximation of the viability kernel for high-dimensional linear sampled-data systems," in *Hybrid Systems: Computation and Control*, Berlin, Germany, 2014, pp. 173–182.

[2] J.-P. Aubin, A. M. Bayen, and P. Saint-Pierre, *Viability Theory: New Directions*, 2nd ed. Springer Verlag, 2011.

[3] F. Blanchini and S. Miani, *Set-Theoretic Methods in Control*. Springer, 2008.

[4] G. Goodwin, J. Aguero, M. Cea Garridos, M. Salgado, and J. Yuz, "Sampling and sampled-data models: The interface between the continuous world and digital algorithms," *IEEE Control Systems Magazine*, vol. 33, no. 5, pp. 34–53, 2013.

[5] L. Magni, D. M. Raimondo, L. Bossi, C. D. Man, G. D. Nicolao, B. Kovatchev, and C. Cobelli, "Model predictive control of type 1 diabetes: An in silico trial," *Journal of Diabetes Science and Technology*, vol. 1, no. 6, pp. 804–812, 2007.

[6] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin, "A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games," *IEEE Transactions on Automatic Control*, vol. 50, no. 7, pp. 947–957, July 2005.

[7] I. M. Mitchell, S. Kaynama, M. Chen, and M. Oishi, "Safety preserving control synthesis for sampled data systems," *Nonlinear Analysis: Hybrid Systems*, vol. 10, pp. 63–82, 2013.

[8] P. Saint-Pierre, "Approximation of the viability kernel," *Applied Mathematics and Optimization*, vol. 29, no. 2, pp. 187–209, Mar 1994.

[9] S. Kaynama and M. Oishi, "A modified Riccati transformation for decentralized computation of the viability kernel under LTI dynamics," *IEEE Transactions on Automatic Control*, vol. 58, no. 11, pp. 2878–2892, 2013.

[10] I. M. Mitchell, "Scalable calculation of reach sets and tubes for nonlinear systems with terminal integrators: A mixed implicit explicit formulation," in *Hybrid Systems: Computation and Control*, Chicago, IL, 2011, pp. 103–112.

[11] P.-A. Coquelin, S. Martin, and R. Munos, "A dynamic programming approach to viability problems," in *IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*, 2007, pp. 178–184.

[12] D. D. Bremner, "On the complexity of vertex and facet enumeration for convex polytopes," Ph.D. dissertation, Montreal, QC, Canada, 1998.

[13] W.-H. Chen, D. J. Ballance, and J. O'Reilly, "Optimisation of attraction domains of nonlinear MPC via LMI methods," in *American Control Conference*, Arlington, VA, 2001, pp. 3067–3072.

[14] A. Alessio, M. Lazar, A. Bemporad, and W. Heemels, "Squaring the circle: An algorithm for generating polyhedral invariant sets from ellipsoidal ones," *Automatica*, vol. 43, no. 12, pp. 2096–2103, 2007.

[15] J. Maidens, S. Kaynama, I. M. Mitchell, M. Oishi, and G. A. Dumont, "Lagrangian methods for approximating the viability kernel in high-dimensional systems," *Automatica*, vol. 49, no. 7, pp. 2017–2029, 2013.

[16] S. Kaynama, J. Maidens, M. Oishi, I. M. Mitchell, and G. A. Dumont, "Computing the viability kernel using maximal reachable sets," in *Hybrid Systems: Computation and Control*, Beijing, 2012, pp. 55–63.

[17] C. Le Guernic and A. Girard, "Reachability analysis of linear systems using support functions," *Nonlinear Analysis: Hybrid Systems*, vol. 4, no. 2, pp. 250–262, 2010.

[18] G. Frehse, C. Le Guernic, A. Donze, S. Cotton, R. Ray, O. Lebeltel, R. Ripado, A. Girard, T. Dang, and O. Maler, "SpaceEx: Scalable verification of hybrid systems," in *23rd International Conference on Computer Aided Verification*, G. Gopalakrishnan and S. Qadeer, Eds. Springer, 2011, pp. 1–16.

[19] A. B. Kurzhanski and P. Varaiya, "Ellipsoidal techniques for reachability analysis: internal approximation," *Systems Control Letters*, vol. 41, pp. 201–211, 2000.

[20] A. Majumdar and R. Tedrake, "Robust online motion planning with regions of finite time invariance," in *Algorithmic Foundations of Robotics*, E. Frazzoli, T. Lozano-Perez, N. Roy, and D. Rus, Eds. Springer Berlin Heidelberg, 2013, pp. 543–558.

[21] M. Tobenkin, I. Manchester, and R. Tedrake, "Invariant funnels around trajectories using sum-of-squares programming," in *IFAC World Congress*, vol. 18, Milano, Italy, 2011, pp. 9218–9223.

[22] R. Tedrake, I. R. Manchester, M. Tobenkin, and J. W. Roberts, "LQR-trees: Feedback motion planning via sums-of-squares verification," *International Journal of Robotics Research*, vol. 29, no. 8, pp. 1038–1052, 2010.

[23] W. Tan and A. Packard, "Searching for control lyapunov functions using sums of squares programming," in *Annual Allerton Conference*, 2004, pp. 210–219.

[24] S. Prajna and A. Jadbabaie, "Safety verification of hybrid systems using barrier certificates," in *Hybrid Systems: Computation and Control*, 2004, pp. 477–492.

[25] D. Henrion and M. Korda, "Convex computation of the region of attraction of polynomial control systems," *preprint arXiv:1208.1751*, 2012.

[26] A. Majumdar, R. Vasudevan, M. M. Tobenkin, and R. Tedrake, "Technical report: Convex optimization of nonlinear feedback controllers via occupation measures," *preprint arXiv:1305.7484*, 2013.

[27] C. Moler and C. Van Loan, "Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later," *SIAM Review*, vol. 45, no. 1, pp. 3–49, 2003.

[28] N. J. Higham, "The scaling and squaring method for the matrix exponential revisited," *SIAM Journal of Matrix Analysis and Applications*, vol. 26, no. 4, pp. 1179–1193, 2005.

[29] C. Standish, "Truncated Taylor series approximation to the state transition matrix of a continuous parameter finite Markov chain," *Linear Algebra and its Applications*, vol. 12, no. 2, pp. 179–183, 1975.

[30] M. Liou, "A novel method of evaluating transient response," *Proceedings of the IEEE*, vol. 54, no. 1, pp. 20–23, 1966.

[31] I. Bárány and Z. Füredi, "Computing the volume is difficult," *Discrete Computational Geometry*, vol. 2, no. 1, pp. 319–326, 1987.

[32] M. Dyer, "Computing the volume of convex bodies: a case where randomness provably helps," *Probabilistic Combinatorics and its Applications*, vol. 44, pp. 123–170, 1991.

[33] C. Schütt and E. Werner, "Polytopes with vertices chosen randomly from the boundary of a convex body," in *Geometric Aspects of Functional Analysis*, ser. LNM 1807, V. Milman and G. Schechtman, Eds. Springer Berlin, 2003, pp. 241–422.

[34] J. Löfberg, "YALMIP: a toolbox for modeling and optimization in MATLAB," in *Computer Aided Control System Design*, 2004, pp. 284–289.

[35] M. Kvasnica, P. Grieder, M. Baotić, and M. Morari, "Multi-Parametric Toolbox (MPT)," in *Hybrid Systems: Computation and Control, LNCS 2993*, R. Alur and G. J. Pappas, Eds. Berlin, Germany: Springer, 2004, pp. 448–462.

[36] K. V. Mardia and P. E. Jupp, *Directional Statistics*, 2nd ed. John Wiley & Sons, Inc, 2000.

[37] M. E. Dyer and A. M. Frieze, "On the complexity of computing the volume of a polyhedron," *SIAM Journal of Computing*, vol. 17, no. 5, pp. 967–974, 1988.

[38] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.

[39] S. Kaynama and C. J. Tomlin, "Benchmark: Flight envelope protection in autonomous quadrotors," in *International Workshop on Applied Verification for Continuous and Hybrid Systems, Part of CPSWeek*, Berlin, Germany, April 2014, http://cps-vo.org/group/ARCH.

[40] I. Cowling, O. Yakimenko, J. Whidborne, and A. Cooke, "Direct method based control system for an autonomous quadrotor," *Journal of Intelligent & Robotic Systems*, vol. 60, pp. 285–316, 2010.

[41] L. Chisci, J. Rossiter, and G. Zappa, "Systems with persistent disturbances: predictive control with restricted constraints," *Automatica*, vol. 37, pp. 1019–1028, 2001.

[42] M. Ghomi, "Optimal Smoothing for Convex Polytopes," *Bulletin of the London Mathematical Society*, vol. 36, no. 4, pp. 483–492, July 2004.

[43] G. E. Bredon, *Topology and Geometry*. Springer, 1993, vol. 139.

[44] J. Pitman, *Probability*. New York: Springer-Verlag, 1993.